



HIGH-AVAILABILITY

Building Reliable, Scalable AR System Solutions

Introduction

This paper will discuss the products, tools and strategies available for building reliable and scalable Action Request System (AR System) solutions. This paper focuses on the real world design and implementation issues related to software and hardware that can be used to enhance the end user experience.

True reliability can only be achieved when there are no Single Points Of Failure (SPOFs) and that usually means having two of everything. However, rather than have double everything, it is usually acceptable to use one component to *backup* two other components, for example a database backup server may also act as a mid-tier backup or an AR Server backup. The specifics of the architecture are up to individual customers to decide, this paper only attempts to highlight what may be achievable.

The Remedy Building Blocks

AR System is comprised of a number of basic building blocks from Remedy (BMC) and 3rd parties that together make a working solution.

Clients

Clients can be traditional AR System applications which use Remedy's proprietary protocols, known as 'native clients', and are usually run on a Windows platform but may include applications built using AR System APIs by 3rd parties.



Figure 1.
Client Types

Web based clients are also supported using IETF standards (web/http) but these do not communicate directly with the AR Server, instead they contact a mid-tier transformation layer. Figure 1. shows how the various clients will be shown in further diagrams in this paper.

Mid-Tier Application Server

Opening access to a wider community of web enabled users the mid-tier component from Remedy runs on a web server and translates traditional AR System forms etc. into html for use by web enabled clients. The mid-tier server then converts responses from the clients back into traditional AR System responses.

A mid-tier server can be implemented as a stand-alone configuration but that will not provide a highly available solution. Load balancing is a better way to provide resilience and also provides a scalable framework. Clustering the mid-tier in addition to load balancing will provide good recovery and improved availability as well as potentially reducing the investment in hardware, as simple networking components can be used.




Figure 2.
Web Switch
and Mid-Tier

Multiple Mid-Tier (MT) application servers can be implemented in parallel, as shown in Figure 2., but the clients must continue to communicate with only one server as it is the mid-tier server that maintains state. This can be achieved either with a static form of load sharing (where groups of users use a specific MT server) or where an active web load balancer maintains *sticky* connections for each client-server relationship.

The mid-tier web servers are a relatively new offering from Remedy and this 'layer' continues to evolve.

AR System Server Tier


The core intelligence of an AR System is held in a number of server processes that deal with client communications, Email, escalations, etc. The introduction of 'Server Groups' was made with AR System 6.x, providing AR System servers with the capability to offer some basic monitoring and overlap capability. Typically client communications can be supported on all nodes within a group but many of the functions that an AR System Server Group must perform will only be operated on one node at any given time.

To assist with load distribution, rather than balancing, additional servers can be deployed to perform tasks like reports or Distributed Server Option (DSO) operations.



The AR System (6.x+) components have a basic form of inbuilt resilience but have no ability to re-start components that have failed and require each node to have fully duplicated licenses, which can be expensive!

Figure 3.
AR Server Tier



The best performance from an AR System server is achieved when the primary components are on the same node that is also running the database. Typically customers deploy reporting and other less critical facilities on other nodes that have a shared access (remote) to the database as shown in Figure 3.

Distributed Server Option

The Distributed Server Option (DSO) is often confused as a way of improving reliability and resilience. There is some benefit in terms of reliability when using a DSO solution but the main function of the DSO option is to move requests and processing responsibility between locations, often in different time-zones. This can only be achieved where the nodes are all active. A remote AR System does not take over the responsibilities of a local server when the local one goes off line should a connection or component fail.

Database Tier

Typically implemented on the same server as the primary AR System server, the database server contains all of the data and work flow structures for the entire system. In short, no database – no AR System.

There are a number of vendor options but Oracle is probably the most popular. Most database vendors offer a number different strategies for improving resilience but there is no perfect solution, particularly considering that the AR System 'likes' to be on the same server as the database.

Figure 4.
Database Server Tier




Figure 4. shows two AR System servers connecting to a common database. This does work but it should be remembered that this is a representation of a logical not a practical implementation. Normally the database will run on one of the servers and the other will remotely access the server, as shown in Figure 8.

Improving Resilience

To improve resilience we must take each *layer* and expand the underlying technology, looking for ways to reduce single points of failure as well as looking for opportunities to spread load. Particularly important is ensuring that each layer has improved probability of making a successful and reliable connection to the next layer.

Improving Client Reliability

There is very little that can be achieved to improve client availability beyond providing each user with multiple machines, although, users with 'native' clients have a more direct route. The native route is more direct and wherever possible client access should be provided using native mechanisms, which will improve reliance.

Client to Mid-Tier Reliability (non-native clients)

Web and mobile devices using non-native communications must connect to the AR System servers through a mid-tier server that transforms the communications between the two systems. Multiple mid-tier servers can be deployed in a parallel configuration but as the mid-tier server maintains the state (session) information for each client the client must continue to connect to the same mid-tier server for the duration of a session.

Client to mid-tier communications can be improved in a number of ways and the appropriate solution really depends on a number of factors but the number of client sessions and duration of each session have an influence on the most cost effective solution.

To substantially improve the mid-tier reliability intelligent load balancers should be deployed in-front of the mid-tier servers, which should have active monitoring components to restart failed servers. The load balancers will maintain the *sticky* nature of each connection while ensuring that each new connection is sent to the server with the least load - although currently only round robin is supported. In the event of a server failure the load balancers will move traffic to another mid-tier server but this will not be transparent to the user. Having an active monitoring component on each mid-tier server ensures that if a server fails it is restarted automatically, which will provide optimal performance to the user. This restart facility is currently only offered by High-Availability.Com, as certified by Remedy under the Alliance Advantage Programme.

Selection of the appropriate load balancing (web switch) device at this level is really a matter for the customer. Remedy typically recommend Cisco's Director and High-Availability.Com would usually recommend Alteon's (Nortel's) [web switch](#) but they all perform the same basic functions just with different performance and price characteristics.

Mid-Tier Reliability

There are two fundamental strategies to improving reliability in the mid-tier, both involve more servers but one relies on external load balancers and requires a Remedy license per server, the other involves clustering and one Remedy license per active mid-tier server. Both clustering and load balancing can be used in combination to great effect.

Mid-tier performance is an issue that can only really be addressed by improving hardware specifications or the number of active servers. The mid-tier is not currently a high performance module and there is believed to be some code optimisation in progress to improve the mid-tier performance.

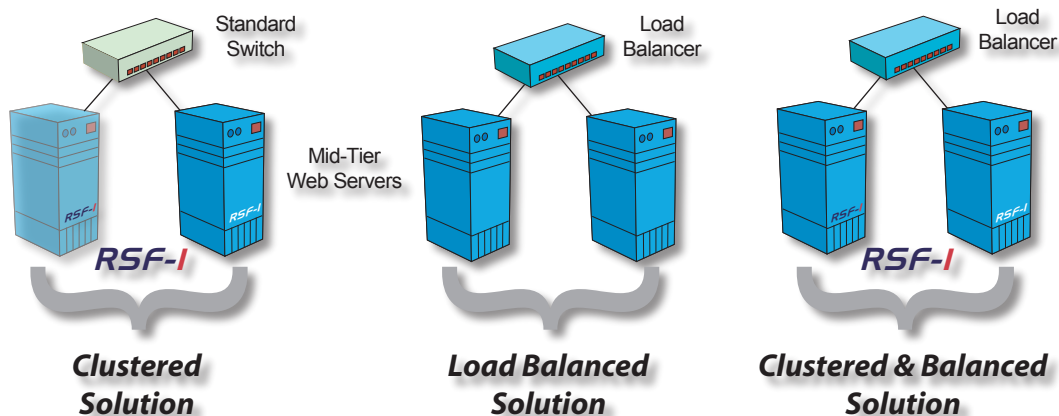


Figure 5. Alternative mid-tier solutions for a high-availability need

Figure 5. shows three example implementations of the mid-tier. The clustered solution requires only one active server and standard networking components. The load balanced solutions require external devices. However, only the clustered solutions restart the mid-tier if it should fail.

Client and Mid-Tier to AR System Reliability

All connections from the mid-tier and native clients use proprietary Remedy protocols but there are a number of areas to look at for improving reliability. At this point the clients (native/mid-tier server) maintain state so even the restart of an AR Server will not be visible to the user if the restart occurs quickly.

Server Groups provide a mechanism to improve the satiability and to a greater extent the resilience of the AR System server tier. Servers can perform some of the tasks in a shared way but the greatest advantage of deploying Server Groups is achieved when splitting the components between the servers as this provides a performance improvement and potentially an improvement in reliability.

Generic load balancers can be deployed with AR System Server Groups to balance load on a round robin basis. Because the AR System protocol is proprietary, the load balancers can not intelligently direct traffic to the appropriate system.



Figure 6.
Client to AR Server
Communications Paths

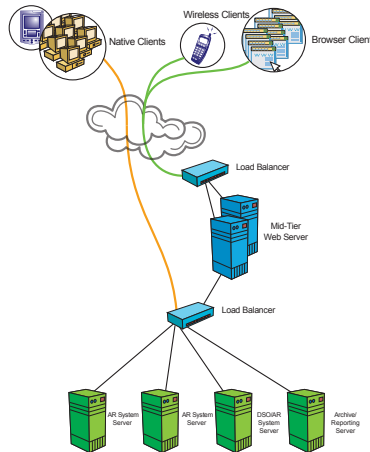


Figure 6. shows the alternative communication paths that clients can take to connect into the AR System server. Note that there are a number of components required, for a resilient structure, that have not been shown here to simplify the diagram but dual (resilient) load balancers, routers and switches are all advisable.

AR System Reliability

The AR System server components are the core infrastructure and are, for performance reasons, normally on the same physical machine as the database, although this is not a requirement and, depending on the machine specifications available, better performance may be achieved using different machines to run the AR System components and the database server.

'Server Groups' provide some improvement in reliability but most of the benefit realised by deploying Server Groups is in the overall system performance. In addition to this, specific servers can be set aside for calculating SLA's or generating reports and handling Email. In the event of a failure another node in the Server Group can take over the functions of the failed node, however, this functionality is limited to the AR System components as well as being limited in the pro-active roll that can be taken. The cost of Server Groups is fairly high as all nodes in the group will need full AR System licenses. Further, some database aware components, like clustering, must also be installed to maintain the accessibility of the underlying database.

The traditional, and arguably the most reliable and cost effective way, to improve reliability at the AR System server *layer* is to use clustering. With clustering you only have one primary active node and you apply hot-back licenses to the alternate node(s). Additional servers can act as reporting servers if required and can be bought into action as a primary server if the primary should fail.

The load balancers that are used in the mid-tier implementation may be *re-used* rather than using additional load balancers. Most of the diagrams in this document show the logical connectivity. However, depending on the site and security requirements it may be desirable to keep the layers separated.

Database Reliability

AR System supports major vendor database products and each has it's own options available to improve reliability. Again clustering is the most common solution to improve reliability at the database *layer*. Replication is a technique that appeals to many organisations but there are performance and data integrity issues that mean replication is normally implemented for disaster recovery or near live reporting rather than a hot-standby solution.

Figure 7.
Alternative database tier
solutions for a high-
availability need

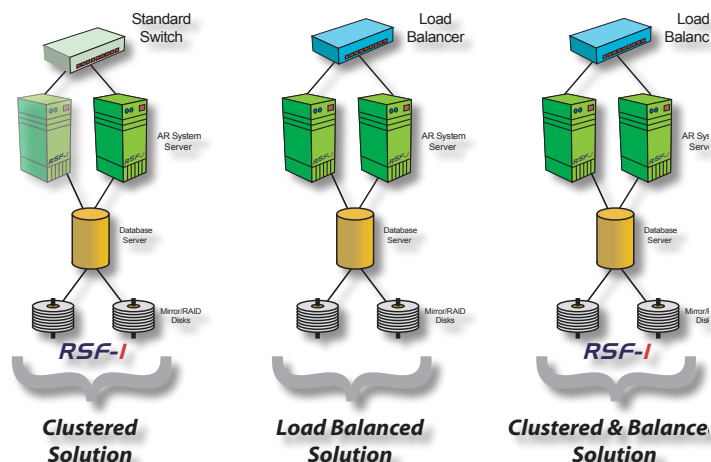


Figure 7. shows the alternative logical deployment methods for a reliable database. The clustered solution is the most simple and most common method and operates by only running one instance of the database and AR System at a given time, this is shown in further detail in Figure 8. The load balanced solutions have multiple instances of AR System server running. The 'load balanced solution' must operate either with a complex and expensive database layer, that will also require an underlying cluster solution like Oracle RAC. The 'clustered & balanced solution' would be our recommended solution; with a single database instance but multiple AR System instances running concurrently.

Hardware Reliability

Removing single points of failure is paramount to providing a reliable solution, no data should be held on one disk spindle alone. Ideally different disks in different cabinets will be used with disk mirroring and/or RAID. Normally Operating System (OS) disk strategies are more than adequate but 3rd party products may be more appropriate in some circumstances. Figure 8. shows a solution that involves disk mirroring between different disk cabinets, that have hardware RAID internally.

The mix of architectures is typical of many installations, where additional compute power is purchased for a primary system but money is saved by accepting a drop in performance when a fault occurs. The example shown is supported but is a rather extreme example, with a Sun E10k failing over to an E450. Normally, on this class of machine, multiple domains would be used to offer a 1st fail over target and another machine as a secondary fail over target.




Figure 8.
Practical Design of Clustered Database

The logical configuration is shown in Figure 8. on the right and the practicable implementation is shown larger on the left. Note that faded 'nodes' indicate that these nodes are not running primary services. These nodes could be used to perform reporting functions etc. while the primary node is running services correctly and reporting functions could be suspended if there is a failure.




Figure 9.
Practical Design of an Active-Active AR System and a Clustered Database

Figure 9. shows a single instance of the database running with multiple instances of the AR System running.



because the load balancing is more *sharing* than *balancing* it is unable to distinguish between the capabilities of the server instances and simply performs a 'round robin' connection system; where each new connection is sent to a server in turn. Note that the secondary system connects to the database on the primary node.

This asymmetric type solution is appropriate for solutions where load is split rather than balanced.

Figure 10.
Summary of Layers
from a Logical View


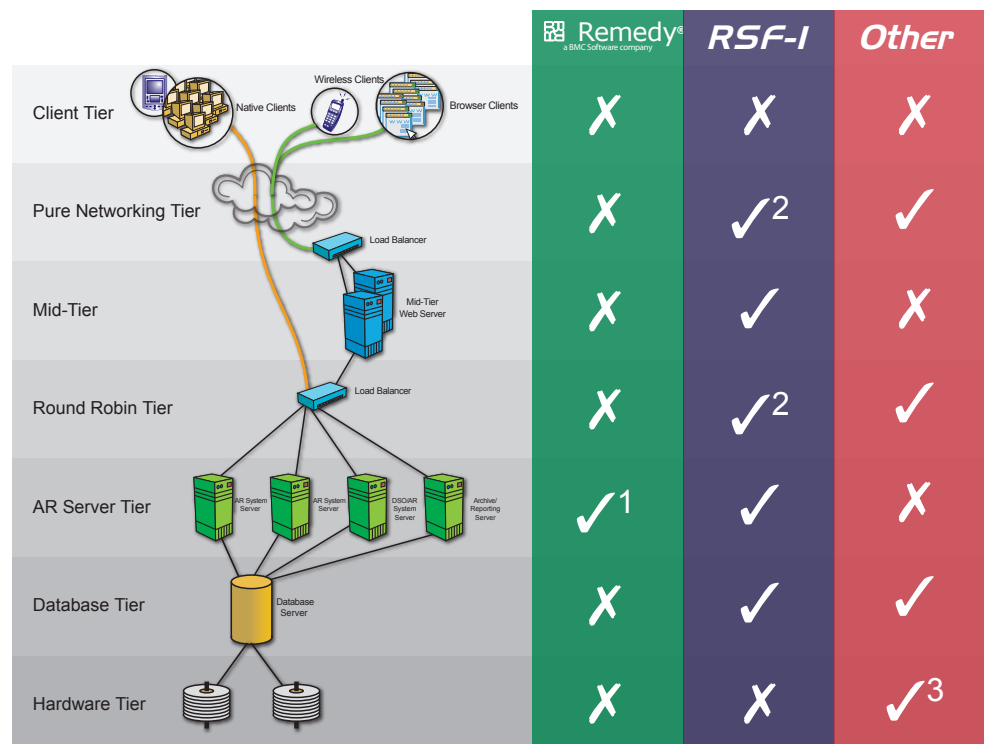


Figure 10. shows an overview of connectivity from a logical point of view.

What Can Improve Reliability At Each 'Layer'

Figure 11.
View of how to address
each reliability for each
'layer'



1 - ONLY 6.x onwards and is limited in capability
2 - Via Virtual IP (VIP) provided by RSF-1
3 - Standard on most OS's or provided by hardware

Figure 11. shows an overview of which products address which needs in each layer.

Cost Implications

The number of variables and specific needs of each customer, not to mention local product pricing policies means that accurate comparisons are very difficult to provide. However, in general terms;

Load balancing switches are expensive and provide limited added value to a AR System solution except for the direction of traffic to mid-tier servers.

For the mid-tier to be built reliably then load balancers and/or clustering must be employed. Both methods add value and can be used usefully together. The benefit of using clustering is that it can restart components as well as 'move' traffic, whereas the load balancers will typically recover (move traffic) ten seconds faster than clustering, but will not restart the failed components.

Server groups require additional AR System licenses which are usually more expensive than clustering, which can be implemented with hot stand-by licenses. Server groups also don't offer automatic restarts of AR System components or any database management, so can only be used to build a highly available solution if clustering has already been employed.

Summary

Remedy's AR System can not be deployed in a highly available way without clustering being deployed within the architecture. While load balancing devices from third parties and 'Server Groups' can offer some resilience and distribution of load, they are neither cheap or comprehensive and depend on other components that without clustering will have single points of failure.

The mid-tier is still maturing and performance is not as good as that of native clients. In addition, sessions will be lost if a mid-tier server fails, where as normally native clients will be able to resume their session transparently on another server. Clustering and load balancing can be used to improve the mid-tier performance and reliability.

Distributed Server Option (DSO) is not targeted or useful for building highly available solutions in general but is focused at moving *data* when everything is working, rather than for duplicating / replicating data to cope with failures.

Clustering is an essential component in building a highly reliable AR System and RSF-1 is only clustering solution which has successfully validated through the Remedy Technology Alliance Programme.

