

CR/COR/001

David Rees  
PA Consulting Group

**GRO**

12 September 1997

Dear David

Below are Pathway's responses to the questions contained in your fax of 10 September. I am faxing this to you today and will speak to you on the telephone to go over any points arising.

1. **On the back of your letter to Peter of the 29<sup>th</sup> August there was a plan which showed a milestone at the 1<sup>st</sup> September of "Scaleability of Escher". Could you please expand on what was behind the milestone and let us know the result?**

This milestone was to take delivery of a version of Riposte which would increase the system limits to cover the 40,000 counters likely to be required. The version of Escher's Riposte system with which we had previously been working had a system limit which restricted it to 5,000 counters. This has now been raised to over 100,000 counters. The release implementing this change was delivered to Pathway on 20<sup>th</sup> August 1997, completing this milestone.

Associated improvements in hand include improving the message replication speed as the number of correspondence servers increases. In addition a multi-clustering technique is about to be tested, and the introduction of sub-groups within Riposte Groups is being considered.

2. **We understand that the overall architecture is 2 Tier Client Server. Are the servers grouped hierarchically, i.e. are there regional servers or is there a set of central servers?**

The architecture is not client-server in the conventional sense. The architecture uses a four-tier system: Counter

Clients, Correspondence Servers, Agent Servers, Host Servers. Counters communicate with Correspondence Servers in a manner which implements the message store in a distributed database manner. Agent Servers communicate with the Correspondence Servers to load messages into Riposte and thence to the Counters, and harvest messages back from the Counters via Riposte. Host servers carry out the processing of incoming and outgoing data.

3. **Are there any plans to change the server hierarchy as the network grows, and if so do you have a timetable for such a change? Would any such change involve software changes in the Riposte kernel?**

No. The architecture outlined above is fixed and does not require the nature or number of hierarchic levels to change as the workload changes. Instead, as the workload grows, all that is necessary is basically to increase the number of units at each level - "sideways expansion". Of course, the network connectivity and bandwidth will need to be kept under review as additional units are added, but this is not an architectural issue.

4. **From our previous discussions our understanding is that Riposte is a very sophisticated messaging software. Will the servers therefore be split between messaging servers and transaction servers?**

Referring to the architectural concepts outlined above, one could say that the Correspondence Servers are messaging servers and the Agent Servers are transaction servers, although this is an approximate description.

5. **Because of the nature of the Riposte system we assume it treats even complex events as messages. Has this led to performance problems when complex operations are performed e.g. starting up a counter, and if so how have these problems been overcome?**

Complex operations can generate a number of messages. These have been taken into account in the sizing model, and as indicated above, as the workload grows the solution is to grow the number of Correspondence Servers and Agent Servers sideways.

Riposte does not operate strictly as a real-time messaging system; except for urgent messages, messages are batched up and sent asynchronously between the Counter and the Campus at a convenient time. Although designed primarily to optimise use of the ISDN links between Outlet and

Campus, this also has the benefit of batching the message load on the Campus.

On the specific area of starting up a counter, there was such an issue which was cleared within Riposte 5.2.

6. **Could you please give an indication of what performance modelling has been carried out or is planned? Did such modelling raise any issues which are still outstanding?**

Detailed workload modelling is in progress and the plans and status are shared with two representatives of the PDA who are specialists in this area and who participate in the activity. The modelling in progress is being done in line with detailed benchmarks of the performance of the systems developed to date, together with extrapolations of their performance to the full expected workload.

There are 10 active issues being progressed as at today's date. These include the need to improve the time taken to harvest the maximum TPS transaction volumes during the overnight TIP batch runs on the Pathway Host systems.

Other modelling activity has concerned the communications infrastructure between the Counters and the Campuses. No problems have been identified in this.

A comprehensive modelling exercise concerning the disk space requirement of the Counter and Correspondence Servers has identified a potential disk space problem at full rollout. To address this, changes have been introduced into Riposte, and will be available at the next release (5.4), to provide message store compression.

7. **Have any other performance reviews, i.e. other than modelling, been carried out? If there have been such reviews could you please indicate the types of issue raised and any that are still outstanding.**

Yes. As described above, detailed benchmarking activities are carried out within The Solution Centre at Bracknell. These have indicated a number of issues relating to the speed of replication of the Riposte message store, and of the creation of indexes into the message store. Improvements introduced by Escher at Release 5.3 have alleviated both of these problems. Benchmarking of the PAS / CMS system was carried out at the end of last year. This identified some areas for improvement and that work has been put in hand.

Yours sincerely

John Dicks  
Director, Customer Requirements