

1 Software Implementation

1.1 Implementation Detail

There are a number of processes involved in calculating and maintaining the relevance and importance values associated with content and entities. This section will outline the various processes that have been implemented to successfully manage BlueSpace Mail Relevance:

- Displaying the calculated BlueSpace Mail Relevance value when browsing the BlueSpace system
- Initial calculation of the BlueSpace Mail Relevance value for each piece of content
- Updating each persons local relevance value on auditable actions
- Updating the local importance values for entities involved in an auditable action
- Updating the global relevance values for each piece of content
- Updating the global importance values for each entity
- Updating the importance values of entities to specific roles
- Updating the local relevance values as a result of changes to local, role or global importance value
- Updating local relevance through manual importance changes

1.1.1 BlueSpace Mail Browsing

Users interacting with the system will be shown relevance values calculated for each message. The value shown to the user is an “effective relevance” based on a combination of the local and global relevance values.

$$\text{Effective relevance} = w_1^x R_k + w_2^x R_g$$

When considering the BlueSpace data model, the local relevance is associated with a user’s mailbox (i.e. pointer to the content), while the global relevance is associated with the actual content. This means that display of the information has minimal impact on the performance of the system.

In the folder views, the user can sort by relevance in the same way they can sort by received date, subject, sender etc.

By combining the global relevance value with the local relevance value, collaborative relevance is obtained. For example, when somebody has been on vacation for a couple of weeks the global relevance value will have changed as a result of how other individuals receiving the same messages have interacted with that message. If each person has read and then deleted that message, the global relevance value will be very low allowing the individual who has been away to concentrate on more relevant messages first.

1.1.2 Initial Calculation

As described in section X: Mathematical Algorithm Framework, the incoming relevance value can be calculated using the following formula:

$$\begin{aligned}
{}^xR^I_k &= \frac{\omega}{W_D} \sum_{j=1}^{n_j} w_i {}^x D_j^i \\
&+ \frac{\upsilon}{W_C} \sum_{j=1}^{n_j} w_j f({}^xC_j^i) \\
&+ \frac{\psi}{W_I} \left((w_1 f({}^g I_s) + w_2 f({}^l I_s) + w_3 \sum_{sv=1}^{n_s} [w_1 f({}^g I_{sv}/n_x) + w_2 f({}^l I_{sv}/n_x)] \right)
\end{aligned}$$

There are three parts to this calculation:

- Calculate the weighted sum of Dichotomous Indicators
- Calculate the weighted sum of the Cardinal Factors
- Calculate the weighted sum of Importance for each entity involved in the message (and message history)

Each of these values is stored separately to facilitate ease of update as a user interacts with the message and impacts the relevance.

The local relevance for each user is calculated when the message is sent/received within the BlueSpace TransAct Server:

Calculate Dichotomous Indicator Weight

Calculate Cardinal Factor Weight

For each recipient (including sender)

Calculate Importance Weight

The initial global relevance value is calculated from the mean of each recipient's local relevance value.

For these initial calculations, the dichotomous indicators and cardinal factors can only use the message details since the recipient has not yet performed any auditable actions. As such, these two calculations are very straightforward and only need to be calculated once for each recipient.

The importance weight calculation has multiple levels of complexity that can be configured based on the infrastructure and performance requirements:

- In the simplest form, the importance weight is calculated on the global and local importance of the entities associated with only the message received.
- A more complex version takes into account the importance of entities in the message history.
- A further option allows both of these calculations to include the importance with respect to the user's role.

Unlike the dichotomous indicators and cardinal factors that are based only on the information associated with the new message, the importance calculation has to look up the importance values from the database. As with other areas of the application, data caching is used to improve the performance of these lookups.

1.1.3 Auditable Actions - Local Relevance

A user's interaction with the BlueSpace system is subject to rigorous audit – the introduction of relevance makes use of the audit mechanism to ensure that

the relevance values associated with a message are updated as a result of the user's interaction with that message, whether that is reading a message, forwarding the message to one or more recipients, replying to a message, moving the message to another folder, deleting the message and even restoring a message.

Each action results in a change to a dichotomous indicator or cardinal factor and the user's local relevance value is recalculated and updated in the database.

The action may also result in a change to the local importance values of entities to the user. However, the local importance value of an entity affects the local relevance of many messages and as a result is not handled by this process (see below).

1.1.4 Auditable Actions – Local Importance

As discussed in section X: Mathematical Algorithm Framework, each action will have an effect on the importance values between a user and an entity. A user's interaction with the system not only affects the relevance associated with individual messages, it also has an affect on the importance values associated with the entities in the system.

This process deals with changes to the local importance values associated between a user and an entity resulting from the auditable actions from a user interacting with content in the system.

The impact of changes to the importance values are far reaching and requires many relevance values to be recalculated. Fortunately, it would be surprising for the local importance values to fluctuate significantly and would expect them to change slowly over a period of time. As such, the mechanism used to recalculate the local relevance scores as a result of changes in importance is handled separately in an asynchronous background process.

1.1.5 Global Relevance

Each user's interaction with the system impacts the relevance placed on the messages they interact with. The global relevance value is used to provide a collaborative element to each user's effective relevance value.

Rather than updating the global relevance value each time a user interacts with the system, which could cause database locking problems, the global relevance value is updated in an asynchronous background process that cycles through all the messages in the system.

```
For each TransMail
```

```
Global Relevance = mean(each user's local relevance)
```

This background process is tuned to spend more time recalculating new and highly relevant messages.

Each calculation is required to access the database to retrieve the local relevance values for each user. The only database locking required is to update a global relevance value. This update is very fast and has no other dependencies ensuring that it does not cause locking issues in the main application.

1.1.6 Global Importance

This process is very similar to the global relevance process but updates the global importance value for each entity. As with global relevance the calculation is straightforward being the mean value of each user's local importance for an entity.

1.1.7 Role Importance

Role importance is an extension to the global importance process that is used to calculate the importance of an entity to each user sharing the same role. Role importance is an optional component that allows the relevance values to take account of the importance of content to a specific role or post. If a user is moved to another role, the relevance of existing content may change dramatically. Role importance allows the relevance values to be weighted towards the role.

1.1.8 Local Relevance

Changes to local and global importance need to be propagated through to the local relevance values associated with specific content.

It would require extensive processing power to propagate every change to importance through to the relevance of all associated content. As such, this background process cycles through all content to recalculate the local relevance values for each user. The recalculation also takes into account the affect of time on the relevance associated with content allowing the relevance to decay over time.

This is an extensive process and could take a significant amount of time (and processing power) to revisit all messages. As with the Global Relevance process, this process is tuned to spend more time processing new and highly relevant messages.

1.1.9 Manual Importance Changes

User feedback on the importance of individuals is likely to have a more direct impact on relevance than a user's "typical" interaction with the system.

To handle these manual changes, this asynchronous process is configured to listen for "manual importance change" events. When such an event is received, each relevance value associated with the changed entity is recalculated.

1.1.10 Deployment Options

The background processes discussed in this section are all processor and time intensive. However, there are a number of deployment options that ensure the ongoing operation of the BlueSpace TransAct Server are not impacted by these processes.

In particular, these background processes can be deployed on additional servers or configured to only run at specific times of day in non-peak periods (for example, overnight) ensuring that end-user performance is not affected by these processes.

Where the user population is high and 24 hour operation required, additional hardware is definitely the best option.

EXHIBIT 2057

metadata slip sheet

BegDoc	REV0001068
BegAttach	REV0001068
File Name	SoftwareImplementation.doc
LastCreateOn	8/18/2005 6:47 AM
LastModifiedDate	8/18/2005 12:15 PM
Custodian	Marston, Paul
Author	Jon Fernyhough
SortDate	8/18/2005 12:15 PM
Office-Revision Number	8
Last Saved By	Jon Fernyhough
Exhibit No	