Subject: Changes in hashing
Posted by mirek on Mon, 01 Jun 2020 13:43:45 GMT

View Forum Message <> Reply to Message

Up until now, hash codes in U++ very strictly 32bit dword. It turns out that using 64bit hash codes on 64bit CPUs is actually faster (more bytes can be processed at once), so I have introduced new type, hash_t, which is 32 bit with 32 bit CPU and 64 bit otherwise and changed the code to compute/use 64 bit hashes instead. Results in about 5% improvement in idmap benchmark...

Practical consideration for user types: If type supports hashing by dword GetHashValue() method, it will continue to work just fine, but might be improved by converting that to hash_t. Template specialisation GetHashValue needs to change the return type hash_t (compiler issues error if it is not).

Mirek