208 have been specified, they are packaged into the knowledge elements 106 in the manner described above with respect to FIG. 2. This occurs at operation 310. It should be appreciated that other data structures may be utilized to package the metadata.

[0037] Once the knowledge elements 106 have been completely assembled, the routine 300 continues to operation 312, where the re-usable software component 104 and the knowledge elements 106 are published. As discussed above, publication may entail storing the knowledge elements 106 and the re-usable software component 104 in a database or another location where they may be discovered and utilized by other software components and re-users. Additional details regarding one discovery and implementation process is provided below with respect to FIG. 4. From operation 312, the routine 300 continues to operation 314, where it ends.

[0038] Referring now to FIG. 4, an illustrative method for discovering and integrating the re-usable software component 104 will be described. The routine 400 begins at operation 402, where a discovery process is performed to locate the re-usable software component 104. As discussed above, in one implementation the re-usable software component 104 and the corresponding knowledge elements 106 are placed in a searchable location, such as a database. In order to discover an appropriate re-usable software component 104, a search is performed of the knowledge elements 106 to identify a suitable re-usable software component 104. For instance, according to implementations a search is made of the capability metadata 202 to identify a re-usable software component 104 with desired capabilities. A search may similarly be made of the affinity metadata 204 and the re-user metadata 210 to locate an appropriate re-usable software component 104 for a particular application.

[0039] Once an appropriate re-usable software component 104 has been discovered, the routine 400 continues to operation 404, where a determination is made as to whether the discovered re-usable software component is to be integrated (also referred to herein as "composited") in an automated fashion or whether a semi-automatic composition is to be performed utilizing input from a re-user of the re-usable software component 104. A fully automated composition does not require any involvement from the re-user. The determination as to whether a fully automated or semi-automatic composition is to be performed may be based upon a user preference, upon data contained in the knowledge elements 106, or upon other factors.

[0040] If a fully automated integration is to be performed, the routine 400 proceeds from operation 404 to operation 406. At operation 406, the re-usable software component 104 is integrated with another software component, such as the application 102 shown in FIG. 1, using the contents of the knowledge elements 106. In particular, the capability metadata 202, affinity metadata 204, derivation metadata 206, and the transformation metadata 208 are utilized to determine how the re-usable software component 104 and the other software component are to be integrated. For instance, the transformation metadata 208 may specify that certain transformation operations, such as union operations or merging operations, are to be performed upon the re-usable software component 104 in order to integrate the two software components. Other types of transformations may also be performed.

[0041] If, at operation 404, it is determined that a semi-automatic integration is to be performed, the routine 400 proceeds from operation 404 to operation 408. At operation 408, options are formulated for the integration using the contents of the knowledge elements 106. For instance, similar to the process performed at operation 406, the types of transformations and other operations to be performed may be identified. A re-user may then be prompted for additional integration options at operation 410. For instance, the user may be asked which type of transformations are to be performed when the transformation metadata 208 specifies that more than one type of transformation may be performed. Once the re-user has specified the integration options, the routine 400 proceeds from operation 410 to operation 412, where the re-usable software component 104 is integrated with the other software component using the contents of the knowledge elements 106 and the user-specified integration options.

[0042] From operations 406 and 412, the routine 400 proceeds to operation 414, where the re-user is permitted to further specify the metadata originally published in the knowledge elements 106. The re-user may also be permitted to add or edit the re-user metadata 210 described above with respect to FIG. 2. In one implementation, the re-user metadata 210 specified by a re-user may be automatically published for use by other re-users. Once this process has completed, the routine 400 proceeds to operation 416, where it ends.

[0043] FIG. 5 shows an illustrative computer architecture for a computer 500 capable of executing the software components described herein for providing, discovering, and integrating a self-describing re-usable software component in the manner presented above. The computer architecture shown in FIG. 5 illustrates a conventional desktop, laptop, or server computer and may be utilized to execute any aspects of the software components presented herein. The computer architecture shown in FIG. 5 includes a central processing unit 502 ("CPU"), a system memory 508, including a random access memory 514 ("RAM") and a read-only memory ("ROM") 516, and a system bus 504 that couples the memory to the CPU 502. A basic input/output system containing the basic routines that help to transfer information between elements within the computer 500, such as during startup, is stored in the ROM 516. The computer 500 further includes a mass storage device 510 for storing an operating system 518, application programs, and other program modules, which have been described in greater detail herein.

[0044] The mass storage device 510 is connected to the CPU 502 through a mass storage controller (not shown) connected to the bus 504. The mass storage device 510 and its associated computer-readable media provide non-volatile storage for the computer 500. Although the description of computer-readable media contained herein refers to a mass storage device, such as a hard disk or CD-ROM drive, it should be appreciated by those skilled in the art that computer-readable media can be any available computer storage media that can be accessed by the computer 500.

[0045] By way of example, and not limitation, computer-readable media may include volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. For example, computer-readable media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid state memory technology, CD-ROM, digital versatile disks ("DVD"), HD-DVD, BLU-RAY, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or