

(51) International Patent Classification:
G06F 15/16 (2006.01)(21) International Application Number:
PCT/US2013/054242(22) International Filing Date:
9 August 2013 (09.08.2013)

(25) Filing Language: English

(26) Publication Language: English


(30) Priority Data:
13/600,309 31 August 2012 (31.08.2012) US(71) Applicant: INTEL CORPORATION [US/US]; 2200
Mission College Boulevard, Santa Clara, California 95054
(US).(72) Inventor: HIREMANE, Radhakrishna; 16924 NW Mesa
View Lane, Beaverton, Oregon 97006 (US).(74) Agent: TROP, Timothy N.; Trop, Pruner & Hu, P.C.,
1616 S. Voss Rd., Ste. 750, Houston, Texas 77057-2631
(US).(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM,
ZW.(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).**Declarations under Rule 4.17:**

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))

(54) Title: ENABLING A CLOUD TO EFFECTIVELY ASSIGN WORKLOADS TO SERVERS



(57) Abstract: In accordance with some embodiments, a public infrastructure as a service (IaaS) user can provide a file, to a cloud service provider, with information about the specific instructions and opcodes that may be used in an application run on the cloud service provider's system. This information may be developed at compile time by the user before the user deploys the workload onto the public IaaS cloud. Thus the user has complete control over the information that is provided.

ENABLING A CLOUD TO EFFECTIVELY ASSIGN WORKLOADS TO SERVERS

Background

[0001] This relates generally to cloud computing.

[0002] In cloud computing, computing and storage capacity is provided to a variety of users based on economy of scale and sharing of resources. This enables users to have effectively limitless capacity without incurring the cost of setting up a data center.

[0003] Infrastructure as a Service or IaaS offers computers, either in physical form or as virtual machines, raw storage, firewalls, load balancers and networks to a variety of users. In public IaaS, the resources are available to any user. Generally, an IaaS cloud provider has a large pool of servers to service users on demand. The cloud user installs its own operating system images and application software and is responsible for their maintenance. The cloud user pays on a per use basis for the service actually used.

[0004] In public IaaS systems, the service provider who provides the cloud for use by all comers, is burdened with the problem that it does not know anything about the nature of the workloads it receives from these users. Generally the reason for this is that privacy and security concerns limit the information that the cloud service provider is provided. The cloud service provider cannot look at the customer's instance on the cloud service provider's servers to understand the application characteristics. Nor can that service provider monitor the physical resource consumption of the customer's applications extensively enough to determine the characteristics of the customer's applications as that would impact the customer application performance.

Brief Description Of The Drawings

[0005] Some embodiments are described with respect to the following figures:

Figure 1 is a schematic depiction for one embodiment to the present invention;

Figure 2 is a flowchart for a sequence according to one embodiment to the present invention; and

Figure 3 is a flowchart for still another embodiment to the present invention.

Detailed Description

[0006] In accordance with some embodiments, a public Infrastructure as a Service (IaaS) user can provide a file, to a cloud service provider, with information about the specific instructions and opcodes that are used in an application to be run on the cloud service provider's system. This information may be developed at compile time by the user before the user deploys the workload onto the public IaaS cloud. Thus the user has complete control over the information that is provided in some embodiments.

[0007] For example, specific instructions and/or opcodes may be known to benefit from special handling on the cloud. Specific types of instructions may benefit from execution on servers with specific capabilities. If the cloud service provider knew that these instructions were present in an application, the cloud service provider could allocate the workload to the most efficient servers for that particular workload.

[0008] In particular, a cloud service provider may increase the consolidation ratios and enhance the efficiency of its infrastructure by packing as many instances as possible in a manner to maximize the individual server utilization. In addition a cloud service provider may avoid conflicting workloads from being placed on the same server. For example, instead of putting two memory bandwidth intensive workloads together on the same server, the workloads may be scheduled on different servers for better management of resources.

[0009] In some embodiments, at compile or deployment time of the application or the virtual machine instance, key application characteristics are captured in a file that may be called catalog. This catalog can be provided by the cloud user to the IaaS cloud service provider when the instance is deployed. Then the cloud service provider has knowledge, about the application characteristics, that can be used for

efficiently scheduling and better mapping workloads to available server resources without needing to monitor or scrutinize the customer's application.

[0010] The catalog need not necessarily contain a runtime profile. For example when encryption is used in an application to be run on the cloud, such as Hypertext Transfer Protocol Secure (HTTPS) or Secure Sockets Layer (SSL), the compiler logs in the application catalog that the application would benefit from Advanced Encryption Standard (AES)-New Instruction (NI). Then the cloud service provider in turn utilizes this knowledge to schedule the workload on an AES-NI enabled server (rather than on an older system without encryption acceleration).

[0011] As another example when expensive floating point is used in the application to be run on the cloud, the catalog information may be used by the cloud service provider to deploy the workload onto servers that have many integrated core (MIC) assists.

[0012] And, as still another example, if the cloud service provider becomes aware that an application is combined with Advanced Vector Extensions (AVX) instruction set extension, as a result of analyzing the catalog, that service provider can match the workload to an available server and increase infrastructure and efficiency.

[0013] In some embodiments, the catalog may contain language specific hints such as that the workload in the virtual machine instance is Java-based and thereby the cloud service provider can schedule the workload on a Java Virtual Machine (JVM) optimized pool of servers.

[0014] Referring to Figure 1, an overview of the system or data center 10 is shown. At the top, the compilation of the application at the cloud user's own site is shown at 1. The application (app) 12 is provided to a compiler 14 which generates the compiled object 16 as well as the catalog 18. As indicated in Figure 1 the catalog 18 may identify that opcodes or instructions, such as AES or AVX instructions, are present in the application that would benefit from being assigned to run on servers with special capabilities.

[0015] Then the customer deploys (#2) the workload into the public IaaS cloud on virtual machine (VM) 20. This means that the cloud receives both the compiled object 16 and the catalog 18. The cloud operating environment 22 with storage 25 includes a catalog reader 23 that reads the catalog information in the virtual machine as indicated by the arrow to the card reader 23. Then the cloud deploys the application to the IaaS hardware best suited to handle the workload to the extent feasible. This deployment may be limited by the availability of servers. In some cases a workload may be delayed until an appropriate server is available. For example if AES instructions are provided, the workload may be provided (#3) to servers with AES-NI and/or AVX capability as indicated at 28 rather than to servers 26 within structure 24 that do not include these capabilities.

[0016] In some embodiments, the compiler develops the catalog at compile time. Then when the cloud customer packages the application in the instance bought from the cloud service provider, the catalog is included with that instance. The catalog reader and the operating environment 22 may be embedded within the cloud operating environment. When a customer deploys the instance, the instance is batched into the cloud operating environment for scheduling on the available IaaS infrastructure 24. The cloud operating environment utilizing the embedded catalog reader understands the characteristics of the application and therefore deploys the workload based on the best infrastructure match.

[0017] Applications with certain capabilities such as AVX or AES-NI capabilities use very specific instructions and opcodes. So a parsing tool, as part of the compiler, locates those instructions and can flag that workload as it is being compiled. For example, the presence of instructions, such as AESENC, AESENCLAST, AESDEC, AESDECLAST, AESKEYGENASSIST, AESIMC, PCLMULQDQ, or their opcodes suggests the application could be compiled to take advantage of AES-NI instructions. The invention is not limited to these opcodes and instructions. Any instruction set instruction that can be identified by its opcode or functionality can be utilized by a compiler to achieve similar objectives.

[0018] In Figure 2 a compiler sequence 14 may be implemented in software, firmware and/or hardware. In software and firmware embodiments, computer executable instructions may be stored in one or more non-transitory computer readable media, such as magnetic, semiconductor and/or optical storage.

[0019] Thus referring to Figure 2, the compiler 14 at the customer site receives the application as indicated in 12 and compiles the code as indicated at 30. A parsing tool within the compiler 14 may be used to locate specific instructions or opcodes as indicated in block 32. Basically a file may be provided with the instructions or opcodes to search for and the parsing tool automatically locates those instructions or opcodes within the code being compiled.

[0020] If the parsing tool finds any of these instructions, at diamond 34, it flags the workload and loads the specific instructions and opcodes found as indicated in block 36. Then a profile catalog may be prepared, as indicated in block 38, to be sent with the compiled object to the cloud service provider. In some embodiments a reader may be sent with the catalog.

[0021] The catalog may be in a particular form or template, provided by the cloud service provider in some embodiments. Also, particular codes may be established to indicate the presence of particular instructions or opcodes in the application. In addition, the catalog may include requests for handling on servers with particular characteristics. The catalog may also indicate how time sensitive the application is so that the cloud can assess whether to delay execution pending availability of a beneficial server. Similarly, the catalog may identify potentially conflicting workloads that should not be assigned to the same server.

[0022] In Figure 3 a workload assigned sequence 40 may be implemented in software, firmware and/or hardware. In software and firmware embodiments, computer executable instructions may be stored in one or more non-transitory computer readable media, such as magnetic, semiconductor and/or optical storage.

[0023] The sequence 40, shown in Figure 3, may be implemented in the cloud operating environment in one embodiment. At the cloud service provider, a workload

assigner that is part of the cloud operating environment 22 receives the workload as indicated in block 42. The catalog reader within the operating environment reads the catalog information as indicated in block 44. Then the application is deployed to the appropriate server or server pool adapted the workload as indicated in block 46.

[0024] In some embodiments, the cloud operating environment 22 is a server or processor based system that executes instructions stored on computer readable media. Specifically, those instructions may be stored in one or more non-transitory computer readable media such as an optical, magnetic or semiconductor storage. For example in Figure 1, such a storage is indicated at 25. That storage may store the instructions needed to implement the sequence shown in Figure 3.

[0025] Likewise, the compiler 14 may include a storage 15 which stores instructions to implement the sequence shown in Figure 2. In some embodiments, more than one storage may be used to store the set of instructions to implement the sequences of Figures 2 and 3.

[0026] In one embodiment, the catalog may be included in one or more packets including a header that identifies the packet or packets as a catalog. Then the packet or packets may be easily located by the parsing tool at compile time. Other techniques for facilitating the location of the catalog among other data may also be used.

[0027] The sequences described herein begin when the application is not yet loaded on the system such as the cloud that it is supposed to run on. The customer or the cloud provider has placed a workload on the cloud provider's queue. Then the cloud provider can pick it up and run it on one of the many servers inside the cloud provider's data center.

[0028] Before placing the workload on a system within the cloud provider's data center, the cloud provider can benefit by knowing if the workload is optimized for special instructions. Then the cloud provider can operate its data center resources more efficiently and customers may get better time to completion in some embodiments.

[0029] For example, a customer may place AVX optimized media workloads on a cloud. For privacy reasons, the cloud service provider may not examine the workload to see that it is AVX optimized or not. The customer has no means to tell the cloud service provider that the workload is AVX optimized or not.

[0030] In some embodiments of the present invention cloud providers and customers can communicate the workload profile in a scalable and automated manner. In this example, if the cloud service provider had AVX enabled servers within its server pool to run the workload, it could deploy the workload on such servers and get better time to completion and handle more workloads over time in some embodiments. Once the workload was placed on an AVX enabled server, the dynamic linked libraries (DLLs) can examine the central processor unit identifier (CPUID) and run the optimized code path. Namely, realizing that the server has AVX capabilities based on its central processing unit identifier, the application can run the optimized code.

[0031] The following clauses and/or examples pertain to further embodiments:

One example embodiment may be a method comprising obtaining information about instructions or opcodes used by an application; and assigning the application to a server of a plurality of servers in a cloud based on that information. The method may also include obtaining information includes obtaining information from a catalog accompanying compiled code for the application. The method may also include said information indicates that the application uses encryption, assigning the application to execute on a server adapted to run encryption instructions. The method may also include determining whether the application uses advanced vector instructions. The method may also include assigning the application to a server adapted to run advanced vector instructions. The method may also include receiving a compiled object together with said information at a cloud for execution. The method may also include reading said information in a cloud operating environment before assigning a workload including the object to a server. The method may also include selecting a server to handle said workload based on said information. The method may also include assigning a server in a public Infrastructure as a Service cloud.

[0032] Another example embodiment may be a non-transitory computer readable media storing instructions that enable a processor to: obtain information about instructions or opcodes used by an application; and assign the application to a server of a plurality of servers in a cloud based on that information. The media may further store instructions to obtain information from a catalog accompanying compiled code for said application. The media may further store instructions to assign the application to execute on a server adapted to run encryption instructions if said information indicates that the application uses encryption. The media may further store instructions to determine whether the application uses advanced vector instructions. The media may further store instructions to assign the application to a server adapted to run advanced vector instructions based on determining that the application uses advanced vector instructions. The media may further store instructions to receive a compiled object together with said information at a cloud for execution. The media may further store instructions to read said information in cloud operating environment before assigning a workload including the object to a server. The media may further store instructions to select a server to handle a workload based on said information. The media may further store instructions to assign a server in a public Infrastructure as a Service cloud.

[0033] One example embodiment may be a non-transitory computer readable media storing instructions to enable a processor to: receive an application; compile code for the application; locate specific instructions or opcodes within said code, and if a specific instruction is found, flag the application; and prepare a profile catalog identifying information about said instructions or opcodes for the application to be read at a cloud service provider's site. The media may include storing instructions to parse the application for encryption instructions. The media may include storing instructions to parse the application to determine whether the application uses advanced vector instructions. The media may include storing instructions to identify said instructions or opcodes at compile time. The media may also include storing instructions to send a compiled object together with said information to a cloud for execution.

[0034] Another example embodiment may be an apparatus comprising a processor to parse instructions or opcodes within an application and if a specific instruction or opcode is found, flag the application and prepare a profile catalog, to identify information about said instruction or opcode for the application, to be read at a cloud service provider's site; and a storage coupled to said processor. The apparatus may include said processor to parse the application for encryption instructions. The apparatus may include said processor to parse the application to determine whether the application uses advanced vector instructions. The apparatus may include said processor to identify said instructions or opcodes at compile time. The apparatus may also include said processor to send a compiled object together with said information to a cloud for execution.

[0035] References throughout this specification to "one embodiment" or "an embodiment" mean that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one implementation encompassed within the present invention. Thus, appearances of the phrase "one embodiment" or "in an embodiment" are not necessarily referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be instituted in other suitable forms other than the particular embodiment illustrated and all such forms may be encompassed within the claims of the present application.

[0036] While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

What is claimed is:

1. A method comprising:
obtaining information about instructions or opcodes used by an application;
and
assigning the application to a server of a plurality of servers in a cloud based on that information.
2. The method of claim 1 wherein obtaining information includes obtaining information from a catalog accompanying compiled code for the application.
3. The method of claim 1 including if said information indicates that the application uses encryption, assigning the application to execute on a server adapted to run encryption instructions.
4. The method of claim 1 including determining whether the application uses advanced vector instructions.
5. The method of claim 4 including assigning the application to a server adapted to run advanced vector instructions.
6. The method of claim 1 including receiving a compiled object together with said information at a cloud for execution.
7. The method of claim 6 including reading said information in a cloud operating environment before assigning a workload including the object to a server.
8. The method of claim 7 including selecting a server to handle said workload based on said information
9. The method of claim 1 including assigning a server in a public Infrastructure as a Service cloud.

10. One or more non-transitory computer readable media storing instructions that enable a processor to:

obtain information about instructions or opcodes used by an application; and
assign the application to a server of a plurality of servers in a cloud based on that information.

11. The media of claim 10 further storing instructions to obtain information from a catalog accompanying compiled code for said application.

12. The media of claim 10 further storing instructions to assign the application to execute on a server adapted to run encryption instructions if said information indicates that the application uses encryption.

13. The media of claim 10 further storing instructions to determine whether the application uses advanced vector instructions.

14. The media of claim 13 further storing instructions to assign the application to a server adapted to run advanced vector instructions based on determining that the application uses advanced vector instructions.

15. The media of claim 11 further storing instructions to receive a compiled object together with said information at a cloud for execution.

16. The media of claim 15 further storing instructions to read said information in cloud operating environment before assigning a workload including the object to a server.

17. The media of claim 16 further storing instructions to select a server to handle a workload based on said information.

18. The media of claim 10 further storing instructions to assign a server in a public Infrastructure as a Service cloud.

19. One or more non-transitory computer readable media storing instructions that enable a processor to:

- receive an application;
- compile code for the application;
- locate specific instructions or opcodes within said code, and if a specific instruction is found, flag the application; and
- prepare a profile catalog identifying information about said instructions or opcodes for the application to be read at a cloud service provider's site.

20. The media of claim 19 further storing instructions to parse the application for encryption instructions.

21. The media of claim 19 further storing instructions to parse the application to determine whether the application uses advanced vector instructions.

22. The media of claim 19 further storing instructions to identify said instructions or opcodes at compile time.

23. The media of claim 19 further storing instructions to send a compiled object together with said information to a cloud for execution.

24. An apparatus comprising:

- a processor to parse instructions or opcodes within an application and if a specific instruction or opcode is found, flag the application and prepare a profile catalog, to identify information about said instruction or opcode for the application, to be read at a cloud service provider's site; and
- a storage coupled to said processor.

25. The apparatus of claim 24 wherein said processor to parse the application for encryption instructions.

26. The apparatus of claim 24, said processor to parse the application to determine whether the application uses advanced vector instructions.
27. The apparatus of claim 24, said processor to identify said instructions or opcodes at compile time.
28. The apparatus of claim 24, said processor to send a compiled object together with said information to a cloud for execution.

1 / 2




FIG. 1

2 / 2




FIG. 2




FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/054242**A. CLASSIFICATION OF SUBJECT MATTER****G06F 15/16(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 15/16; G06F 15/173; G06F 13/00; G06F 21/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: IaaS(Infrastructure as a Service, cloud server, workload, application, compile

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011-0231899 A1 (PULIER ERIC et al.) 22 September 2011 See abstract; paragraphs [6]-[14], [52]-[64], [80]-[92]; claims 1-38 and figures 1, 2A, 2B, 8.	1-2, 6-11, 15-18
A		3-5, 12-14, 19-28
A	JP 2001-005791 A (NEC CORP) 12 January 2001 See abstract; paragraphs [7]-[11]; claims 1-2 and figures 1-2.	1-28
A	US 2010-0106812 A1 (BERNABEU-AUBAN JOSE et al.) 29 April 2010 See abstract; paragraphs [15]-[17], [40]-[47]; claims 1-8 and figure 2.	1-28
A	US 2011-0153824 A1 (CHIKANDO ERIC N. et al.) 23 June 2011 See abstract; paragraphs [10]-[15], [34]-[41]; claims 1-20 and figures 1-3.	1-28

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

28 November 2013 (28.11.2013)

Date of mailing of the international search report

29 November 2013 (29.11.2013)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City,
302-701, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

HONG, Kyoung Ah

Telephone No. +82-42-481-5668



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/054242

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0231899 A1	22/09/2011	WO 2011-091056 A1 WO 2012-100092 A2 WO 2012-100092 A3	28/07/2011 26/07/2012 27/09/2012
JP 2001-005791 A	12/01/2001	None	
US 2010-0106812 A1	29/04/2010	None	
US 2011-0153824 A1	23/06/2011	None	