

1     **METHOD AND SYSTEM FOR DETERMINING COMPATIBILITY OF COMPUTER**  
2     **SYSTEMS**

3     **[0001]**     This application is a continuation-in-part of U.S. Patent Application Nos. 11/535,355  
4     and 11/535,308 filed on September 26, 2006, both of which claim priority from U.S. provisional  
5     patent application no. 60/745,322 filed April 21, 2006, the contents of each being incorporated  
6     herein by reference.

7     **TECHNICAL FIELD**  
8

9     **[0002]**     The present invention relates to information technology infrastructures and has  
10    particular utility in determining compatibility of computer systems in such infrastructures.

11    **BACKGROUND**

12   **[0003]**     As organizations have become more reliant on computers for performing day to day  
13    activities, so to has the reliance on networks and information technology (IT) infrastructures  
14    increased. It is well known that large organizations having offices and other facilities in different  
15    geographical locations utilize centralized computing systems connected locally over local area  
16    networks (LAN) and across the geographical areas through wide-area networks (WAN).

17   **[0004]**     As these organizations grow, the amount of data to be processed and handled by the  
18    centralized computing centers also grows. As a result, the IT infrastructures used by many  
19    organizations have moved away from reliance on centralized computing power and towards  
20    more robust and efficient distributed systems. Distributed systems are decentralized computing  
21    systems that use more than one computer operating in parallel to handle large amounts of data.  
22    Concepts surrounding distributed systems are well known in the art and a complete discussion  
23    can be found in, e.g., "Distributed Systems: Principles and Paradigms"; Tanenbaum Andrew S.;  
24    Prentice Hall; Amsterdam, Netherlands; 2002.

25   **[0005]**     While the benefits of a distributed approach are numerous and well understood, there  
26    has arisen significant practical challenges in managing such systems for optimizing efficiency  
27    and to avoid redundancies and/or under-utilized hardware. In particular, one challenge occurs  
28    due to the sprawl that can occur over time as applications and servers proliferate. Decentralized  
29    control and decision making around capacity, the provisioning of new applications and hardware,

21635070.1

1 and the perception that the cost of adding server hardware is generally inexpensive, have created  
2 environments with far more processing capacity than is required by the organization.

3 [0006] When cost is considered on a server-by-server basis, the additional cost of having  
4 underutilized servers is often not deemed to be troubling. However, when multiple servers in a  
5 large computing environment are underutilized, having too many servers can become a burden.  
6 Moreover, the additional hardware requires separate maintenance considerations; separate  
7 upgrades and requires the incidental attention that should instead be optimized to be more cost  
8 effective for the organization. Heat production and power consumption can also be a concern.  
9 Even considering only the cost of having redundant licenses, removing even a modest number of  
10 servers from a large computing environment can save a significant amount of cost on a yearly  
11 basis.

12 [0007] As a result, organizations have become increasingly concerned with such  
13 redundancies and how they can best achieve consolidation of capacity to reduce operating costs.  
14 The cost-savings objective can be evaluated on the basis of consolidation strategies such as, but  
15 not limited to: virtualization strategies, operating system (OS) level stacking strategies, database  
16 consolidation strategies, application stacking strategies, physical consolidation strategies, and  
17 storage consolidation strategies.

18 [0008] Virtualization involves virtualizing a physical system as a separate guest OS instance  
19 on a host machine. This enables multiple virtualized systems to run on a single physical  
20 machine, e.g. a server. Examples of virtualization technologies include VMware™, Microsoft  
21 Virtual Server™, IBM LPAR™, Solaris Containers™, Zones™, etc.

22 [0009] OS-Level application stacking involves moving the applications and data from one or  
23 more systems to the consolidated system. This can effectively replace multiple operating system  
24 instances with a single OS instance, e.g. system A running application X and system B running  
25 application Y are moved onto system C running application Z such that system C runs  
26 applications X, Y and Z, and system A and B are no longer required. This strategy is applicable  
27 to all operation system types, e.g. Windows™, Linux™, Solaris™, AIX™, HP-UX™, etc.

1 [0010] Database stacking combines one or more database instances at the database server  
2 level, e.g. Oracle™, Microsoft SQL Server™, etc. Database stacking combines data within a  
3 database instance, namely at the table level. Application stacking combines one or more  
4 database instances at the application server level, e.g. J2EE™ application servers, Weblogic™,  
5 WebSphere™, JBoss™, etc.

6 [0011] Physical consolidation moves physical systems at the OS level to multi-system  
7 hardware platforms such as Blade Servers™, Dynamic System Domains™, etc. Storage  
8 consolidation centralizes system storage through storage technologies such as Storage Area  
9 Networks (SAN), Network Attached Storage (NAS), etc.

10 [0012] The consolidation strategies to employ and the systems and applications to be  
11 consolidated are to be considered taking into account the specific environment. Consolidation  
12 strategies should be chosen carefully to achieve the desired cost savings while maintaining or  
13 enhancing the functionality and reliability of the consolidated systems. Moreover, multiple  
14 strategies may often be required to achieve the full benefits of a consolidation initiative.

15 [0013] Complex systems configurations, diverse business requirements, dynamic workloads  
16 and the heterogeneous nature of distributed systems can cause incompatibilities between  
17 systems. These incompatibilities limit the combinations of systems that can be consolidated  
18 successfully. In enterprise computing environments, the virtually infinite number of possible  
19 consolidation permutations which include suboptimal and incompatibility system combinations  
20 make choosing appropriate consolidation solutions difficult, error-prone and time consuming.

21 [0014] It is therefore an object of the following to obviate or mitigate the above-described  
22 disadvantages.

## 23 SUMMARY

24 [0015] In one aspect, a method for determining a consolidation solution for a plurality of  
25 computer systems is provided comprising obtaining a data set comprising a compatibility score  
26 for each pair of the plurality of computer systems, each the compatibility score being indicative  
27 of the compatibility of one of the plurality of computer systems with respect to another of the

21635070.1

1 plurality of computer systems; determining one or more candidate transfer sets each indicating  
2 one or more of the computer systems capable of being transferred to a target computer system;  
3 selecting a desired one of the one or more candidate transfer sets; and providing the desired one  
4 as a consolidation solution.

5 **[0016]** In another aspect, there is provided method for determining compatibilities for a  
6 plurality of computer systems comprising: obtaining at least one transfer set indicating one or  
7 more of the computer systems capable of being transferred to a target computer system;  
8 evaluating compatibilities of the one or more computer systems against the target computer  
9 system to obtain a first compatibility score; evaluating compatibilities of each the one or more of  
10 the computer systems against each other to obtain a second compatibility score; and computing  
11 an overall compatibility score for the transfer set using the first and second compatibility scores.

12 **[0017]** In yet another aspect, there is provided a computer program for determining  
13 compatibilities for a plurality of computer systems comprising an audit engine for obtaining  
14 information pertaining to the compatibility of the plurality of computer systems; an analysis  
15 engine for generating a compatibility score for each pair of the plurality of systems based on the  
16 information that is specific to respective pairs; and a client for displaying the compatibility score  
17 on an interface, the interface being configured to enable a user to specify parameters associated  
18 with a map summarizing the compatibility scores and to define and initiate a transfer of one or  
19 more of the computer systems to a target computer system.

## 20 BRIEF DESCRIPTION OF THE DRAWINGS

21 **[0018]** An embodiment of the invention will now be described by way of example only with  
22 reference to the appended drawings wherein:

23 **[0019]** Figure 1 is a block diagram of an analysis program for evaluating the compatibility of  
24 computer systems to identify consolidation solutions.

25 **[0020]** Figure 2 is a more detailed diagram of the analysis program depicted in Figure 1.

- 1    **[0021]**    Figure 3 is a block diagram illustrating a sample consolidation solution comprised of  
2    multiple transfers.
- 3    **[0022]**    Figure 4 is an example of a rule-based compatibility analysis map.
- 4    **[0023]**    Figure 5 is an example of a workload compatibility analysis map.
- 5    **[0024]**    Figure 6 is an example of an overall compatibility analysis map.
- 6    **[0025]**    Figure 7 is a schematic block diagram of an underlying architecture for implementing  
7    the analysis program of Figure 1.
- 8    **[0026]**    Figure 8 is a table mapping audit data sources with consolidation strategies.
- 9    **[0027]**    Figure 9 is a process flow diagram of the compatibility and consolidation analyses.
- 10   **[0028]**    Figure 10 is a process flow diagram illustrating the loading of system data for  
11   analysis.
- 12   **[0029]**    Figure 11 is a high level process flow diagram for a 1-to-1 compatibility analysis.
- 13   **[0030]**    Figure 12 is a table showing example rule sets.
- 14   **[0031]**    Figure 13 is a table showing example workload types.
- 15   **[0032]**    Figure 14 is a process flow diagram for the 1-to-1 compatibility analysis.
- 16   **[0033]**    Figure 15 is a flow diagram illustrating operation of the rule engine analysis.
- 17   **[0034]**    Figure 16 shows an example rule set.
- 18   **[0035]**    Figure 17 is a flow diagram of the 1-to-1 rule-based compatibility analysis.
- 19   **[0036]**    Figure 18 is a flow diagram illustrating the evaluation of a rule set.
- 20   **[0037]**    Figure 19 is an example of the rule-based compatibility analysis result details.

- 1    **[0038]**    Figure 20 is a flow diagram of workload data extraction process.
- 2    **[0039]**    Figure 21 is a flow diagram of the 1-to-1 workload compatibility analysis.
- 3    **[0040]**    Figure 22 is an example of the workload compatibility analysis result details.
- 4    **[0041]**    Figure 23 is an example of the overall compatibility analysis result details.
- 5    **[0042]**    Figure 24(a) is a high level process flow diagram of the multi-dimensional  
6    compatibility analysis.
- 7    **[0043]**    Figure 24(b) is a flow diagram showing the multi-dimensional analysis.
- 8    **[0044]**    Figure 24(c) is a flow diagram showing use of a rule set in an N-to-1 compatibility  
9    analysis.
- 10   **[0045]**    Figure 24(d) is a flow diagram showing use of a rule set in an N-by-N compatibility  
11   analysis.
- 12   **[0046]**    Figure 25 is a process flow diagram of the multi-dimensional workload compatibility  
13   analysis.
- 14   **[0047]**    Figure 26 is an example multi-dimensional compatibility analysis map.
- 15   **[0048]**    Figure 27 is an example of the multi-dimensional compatibility analysis result details  
16   for a rule set.
- 17   **[0049]**    Figure 28 is an example of the multi-dimensional workload compatibility analysis  
18   result details.
- 19   **[0050]**    Figure 29 is a process flow diagram of the consolidation analysis.
- 20   **[0051]**    Figure 30 is a process flow diagram of an auto fit algorithm used by the consolidation  
21   analysis.

- 1    **[0052]**    Figure 31 is an example of a consolidation solution produced by the consolidation  
2    analysis.
- 3    **[0053]**    Figure 32 shows an example hierarchy of analysis folders and analyses.
- 4    **[0054]**    Figure 33 shows the screen for creating and editing the analysis input parameters.
- 5    **[0055]**    Figure 34 shows an example of the rule set editor screen.
- 6    **[0056]**    Figure 35 shows an example of the screen for editing workload settings.
- 7    **[0057]**    Figure 36 shows an example screen to edit the importance factors used to compute the  
8    overall compatibility score.
- 9    **[0058]**    Figure 37 is an example 1-to-1 compatibility map.
- 10   **[0059]**    Figure 38 is example configuration compatibility analysis details.
- 11   **[0060]**    Figure 39 is an example 1-to-1 compatibility map for business constraints.
- 12   **[0061]**    Figure 40 is an example 1-to-1 workload compatibility map.
- 13   **[0062]**    Figure 41 is an example workload compatibility report.
- 14   **[0063]**    Figure 42 is an example workload details report with stacked workloads.
- 15   **[0064]**    Figure 43 is an example of a 1-to-1 overall compatibility map.
- 16   **[0065]**    Figure 44 is an example of the 1-to-1 overall compatibility details report.
- 17   **[0066]**    Figure 45 shows the workload details of the overall compatibility report.
- 18   **[0067]**    Figure 46 shows example transfers on a compatibility map with net effect off.
- 19   **[0068]**    Figure 47 shows example transfers on a compatibility map with net effect on.
- 20   **[0069]**    Figure 48 is an example multi-dimensional compatibility details report.

1 [0070] Figure 49 shows an example of the consolidation analysis (auto fit) input screen.

2 [0071] Figure 50 is an example overall compatibility map for the consolidation solution.

3 [0072] Figure 51 is an example consolidation summary report.

4 [0073] Figure 52 shows the example transfers that comprise the consolidation solution.

## 5 DETAILED DESCRIPTION OF THE DRAWINGS

### 6 Analysis Program Overview

7 [0074] A block diagram of an analysis program 10 for determining compatibilities in  
8 computing environment 12 is provided in Figure 1. The analysis program 10, accessed through a  
9 computer station 14, gathers data 18 pertaining to a collection of systems to be consolidated 16.  
10 The analysis program 10 uses the gathered data 18 to evaluate the compatibility of the computer  
11 systems 28 and provide a roadmap 20 specifying how the original set of systems can be  
12 consolidated to a smaller number of systems 22.

13 [0075] The following provides an overview of the principles and functionality related to the  
14 analysis program 10 and its environment depicted in Figure 2.

### 15 System Data Parameters

16 [0076] A distinct data set is obtained for each system 16 to contribute to the combined  
17 system data 18 shown in Figure 2. Each data set comprises one or more parameters that relate  
18 preferably to technical 24, business 26 and workload 28 characteristics or features of the  
19 respective system 16. The parameters can be evaluated by scrutinizing program definitions,  
20 properties, objects, instances and any other representation or manifestation of a component,  
21 feature or characteristic of the system 16. In general, a parameter is anything related to the  
22 system 16 that can be evaluated, quantified, measured, compared etc.

23 [0077] Examples of technical parameters relevant of the consolidation analysis include the  
24 operating system, OS version, patches, application settings, hardware devices, etc.



1 [0078] Examples of business parameters of systems relevant to the consolidation analysis  
2 include the physical location, organization department, data segregation requirements, owner,  
3 service level agreements, maintenance windows, hardware lease agreements, software licensing  
4 agreements, etc.

5 [0079] Examples of workload parameters relevant to consolidation analysis include various  
6 resource utilization and capacity metrics related to the system processor, memory, disk storage,  
7 disk I/O throughput and network bandwidth utilization.

## 8 **System and Entity Models**

9 [0080] The system data parameters associated with a system 16 comprise the system model  
10 used in the analyses.

11 [0081] In the following examples, a source system refers to a system from which  
12 applications and/or data are to be moved, and a target server or system is a system to which such  
13 applications and/or data are to be moved. For example, an underutilized environment having two  
14 systems 16 can be consolidated to a target system (one of the systems) by moving applications  
15 and/or data from the source system (the other of the systems) to the target system.

16 [0082] The computer systems 16 may be physical systems, virtual systems or hypothetical  
17 models. In contrast to actual physical systems, hypothetical systems do not currently exist in the  
18 computing environment 12. Hypothetical systems can be defined and included in the analysis to  
19 evaluate various types of “what if” consolidation scenarios. Hypothetical targets can be used to  
20 simulate a case where the proposed consolidation target systems do not exist in the environment  
21 12, e.g. for adding a system 16. Similarly, hypothetical source systems can be used to simulate  
22 the case where a new application is to be introduced into the environment 12 and “forward  
23 consolidated” onto existing target systems 16.

24 [0083] Hypothetical systems can be created through data imports, cloning from actual  
25 systems models, and manual specification by users, etc. The system model can be minimal  
26 (sparse) or include as much data as an actual system model. These system models may also be  
27 further modified to address the analysis requirements.

21635070.1

1 [0084] The compatibility analysis can also be generalized to evaluate entities beyond  
2 physical, virtual or hypothetical systems. For example, entities can be components that comprise  
3 systems such as applications and database instances. By analysing the compatibility of database  
4 instances and database servers with database stacking rule sets, database consolidation can also  
5 be assessed. Similarly, application consolidation can be evaluated by analyzing application  
6 servers and instances with application stacking rules. The entity could also be a logical  
7 application system and technical data can pertain to functional aspects and specifications of the  
8 entity.

9 [0085] It will therefore be appreciated that a “system” or “computer system” hereinafter  
10 referred, can encompass any entity which is capable of being analysed for any type of  
11 compatibility and should not be considered limited to existing or hypothetical physical or virtual  
12 systems etc.

### 13 Consolidation and Transfers

14 [0086] Consolidation as described above can be considered to include one or more  
15 “transfers”. The actual transfer describes the movement of a single source entity onto a target,  
16 wherein the specification identifies the source, target and transfer type. The transfer type (or  
17 consolidation strategy) describes how a source entity is transferred onto a target, e.g.  
18 virtualization, OS stacking etc.

19 [0087] A transfer set 23 (see Figure 3) can be considered one or more transfers that involve a  
20 common target, wherein the set specifies one or more source entities, the target and a transfer  
21 type.

22 [0088] A consolidation solution (or roadmap) is one or more transfer sets 23 based on a  
23 common pool of source and target entities. As can be seen in Figure 2, the consolidation  
24 roadmap can be included in the analysis results 20. Each source or target entity is referenced at  
25 most one time by the transfer sets that comprise the solution.

26 [0089] Figure 3 shows how an example pool 24 of 5 systems (S1, S2, S3, S4 and S5) can be  
27 consolidated through 2 transfer sets 23: stack S1 and S2 onto S3, and stack S4 onto S5. The

21635070.1

transfer sets 23 include 3 transfers, and each system 16 is referenced by the transfer sets 23 only once. In the result, a consolidated pool 26 of 2 systems is achieved.

[0090] It will be appreciated that the principles described herein support many transformation strategies and consolidation is only one example.

## **Compatibility Analyses**

[0091] The following discusses compatibilities between systems 16 based on the parameters to determine if efficiencies can be realized by consolidating either entire systems 16 or aspects or components thereof.

[0092] The analyses employ differential rule sets 28 to evaluate and quantify the compatibility of systems 16 with respect to technical configuration and business related factors comprised in the gathered system data 18. Similarly, workload compatibility of a set of systems 16 is assessed using workload stacking and scoring algorithms 30. The results of configuration, business and workload compatibility analyses are combined to produce an overall compatibility score for a set of systems 16.

[0093] In addition to compatibility scores, the analysis provides details that account for the actual scores. The scores can be presented in color coded maps 32, 34 and 36 that illustrate patterns of the compatibility amongst the analyzed systems as shown in Figures 4, 5 and 6 respectively.

## **Analysis Modes**

[0094] A collection of systems 16 to be consolidated can be analyzed in one of three modes: 1-to-1 compatibility, multi-dimensional compatibility and consolidation analyses. These analyses share many common aspects but can be performed independently.

[0095] The 1-to-1 compatibility analysis evaluates the compatibility of every possible source-target pair combination in the collection of systems 16 on a 1-to-1 basis. This analysis is useful in assessing single transfer consolidation candidates. In practice, it may be prudent to

consolidate systems 16 incrementally and assess the impact of each transfer before proceeding with additional transfers.

[0096] The multi-dimensional compatibility analysis evaluates the compatibility of transfer sets that can involve multiple sources being transferred to a common target. The analysis produces a compatibility score for each specified transfer set 23 by evaluating the compatibility of the systems 16 that comprise the transfer set 23.

[0097] The consolidation analysis searches for a consolidation solution that minimizes the number of remaining source and target entities after the proposed transfers are applied, while meeting requisite compatibility constraints. This analysis employs the multi-dimensional compatibility analysis described above to evaluate the compatibility of postulated transfer sets.

[0098] The analysis program 10 performs consolidation analyses for virtualization and stacking strategies as will be explained in greater detail below, however, it will be appreciated that other consolidation strategies may be performed according to similar principles.

## **Analysis Program Architecture**

[0099] A block diagram of the analysis program 10 is shown in Figure 7. The flow of data 18 through the program 10 begins as an audit engine 40 pulls audit data 18 from audited environments 42. The data works its way up to the web client 44 which displays an output on a user interface, e.g. on computer 14. The program 10 is preferably a client-server application that is accessed via the web client 44.

[00100] An audit engine 40 communicates over one or more communication protocols with audited environments 42 comprised of the actual systems 16 being analysed. The audit engine 40 typically uses data acquisition adapters to communicate directly with the end points (e.g. servers) or through software systems that manage the end points (e.g. management frameworks 46 and/or agent instrumentation 48 and/or direct access 50).

[00101] Alternatively, system data 18 can be imported from third party tools, e.g. inventory applications, performance monitoring tools etc., or can be obtained through user data entry.

21635070.1

1 Examples of such third-party data having file formats that can be imported include comma  
2 separated values (CSV), extensible markup language (XML) and well formatted text files such as  
3 those generated by the UNIX™ system activity reporter (SAR).

4 [00102] The audit engine 40, uses a set of audit request templates 52 that define the data 18 to  
5 acquire from the systems 16. Once collected, the data 18 is stored in an audit data repository 54.  
6 Data 18 referenced by the analysis rule sets 28 and required by the workload analysis 30 are  
7 extracted and stored in separate data caches 56 and 58 respectively.

8 [00103] Aliases 60, differential rule sets 28, workload data types 30 and benchmark  
9 specifications comprise some of the analysis-related metadata 62 definitions used by the program  
10 10. Aliases 60 extract and normalize system data 18 from a variety of data sources to a common  
11 model for analysis. Rule sets 28 examine system compatibility with respect to technical  
12 configuration and business-related factors. The workload definitions 30 specify the system  
13 resource parameters and benchmarks for analyzing workload compatibility.

14 [00104] An analysis engine 64 is also provided, which comprises compatibility and  
15 consolidation analysis engines 66 and 68 respectively. The compatibility analysis evaluates the  
16 compatibility of systems 16 through rule sets 28 and workload stacking algorithms 30. The  
17 consolidation analysis engine 68 leverages the compatibility analysis and employs constraint-  
18 based optimization algorithms to find consolidation solutions that allows the environment 12 to  
19 operate with fewer systems 16.

20 [00105] The program 10 has a report engine 70 that utilizes report templates 72 for generating  
21 reports that convey the analysis results. Typically, the program 10 includes a web interface layer  
22 74 that allows web client 44 users to enter settings, initiate an audit or analysis, view reports etc.

## 23 **Analysis Data Sources**

24 [00106] The audit data 18 can be acquired using tools such as the table 76 shown in Figure 8  
25 that illustrate the various types of configuration settings that are of interest and from which  
26 sources they can be obtained. Figure 8 also provides a mapping to where the sample workload  
27 data can be obtained. In Figure 8, a number of strategies 78 and sub-strategies 80 map to various

21635070.1

1 configuration and workload sources, collectively referred to by numeral 82. As discussed  
2 making reference to Figure 8, the strategies 78 may relate to database consolidation, OS-level  
3 stacking, application server stacking, virtualization, and many others. Each strategy 78 includes  
4 a set of sub-strategies 80, which in turn map to specific rule sets 28. The rule sets 28, which will  
5 be explained in greater detail below, determine whether or not a particular setting or system  
6 criterion/criteria have been met and thus how different one system 16 is to the next. The rule  
7 sets 28 can also indicate the cost of remediating such differences.

8 **[00107]** The table 76 lists the supported consolidation strategies and the relevant data sources  
9 that should be audited to perform the corresponding consolidation analysis. In general,  
10 collecting more basis data 18 improves the analysis results. The table 76 enables the analysis  
11 program 10 to locate the settings and information of interest based on the strategy 78 or sub-  
12 strategy 80 (and in turn the rule set 28) that is to be used to evaluate the systems 16 in the  
13 environment 12. The results can be used to determine source/target candidates for analysing the  
14 environment for the purpose of, e.g. consolidation, compliance measures etc.

## 15 **Analysis Process Overview**

16 **[00108]** Referring now to Figure 9, a process flow diagram illustrates the data flow for  
17 performing the compatibility and consolidation analyses discussed above. The flow diagram  
18 outlines four processes: a data load and extraction process (A), a 1-to-1 compatibility analysis  
19 process (B), a multi-dimensional compatibility analysis process (C), and a consolidation analysis  
20 process (D).

21 **[00109]** In process A, the system data 18 collected via audits or imports as discussed above is  
22 prepared for use by the analyses. The compatibility and consolidation analyses processes B, C  
23 and D can be performed independently. The analyses share a common analysis input  
24 specification and get system data 18 from the data repository 54 and caches 56 and 58. The  
25 multi-dimensional compatibility and consolidation analyses take additional inputs in the form of  
26 a consolidation solution and auto fit input parameters 84 and 86 respectively.

1 [00110] The 1-to-1 compatibility analysis process B evaluates the compatibility of each  
2 system pair on a 1-to-1 basis. In contrast, the multi-dimensional analysis process C evaluates the  
3 compatibility of each transfer set 23 in the consolidation solution that was specified as part of its  
4 input.

5 [00111] The consolidation analysis process D searches for the best consolidation solution that  
6 fulfills the constraints defined by the auto fit input 86. The consolidation analysis employs the  
7 multi-dimensional compatibility analysis C to assess potential transfer set candidates.

## 8 **Data Load and Extraction**

9 [00112] A process flow diagram for the data load and extraction process A is illustrated in  
10 Figure 10. System data including technical configuration, business related and workload  
11 collected through audits, data import and user input are prepared for use by the analyses  
12 processes B, C and D.

13 [00113] When system data 18 and attributes are loaded into the analysis program 10, they are  
14 stored in the audit data repository 54 and system attribute table 55, respectively. As well, system  
15 data 18 referenced by rule set items 28, workload types 30 and benchmarks are extracted and  
16 loaded into their respective caches 56, 58. Alias specifications 60 describe how data can be  
17 extracted and if necessary, normalized from a variety of data sources.

18 [00114] The data repository 54 and caches 56 and 58 thus store audited data 18, system  
19 attributes, the latest rule set data, historical workload data and system workload benchmarks.

## 20 **1-to-1 Compatibility Analysis**

21 [00115] A high level flow diagram of the 1-to-1 compatibility analysis is shown in Figure 11.  
22 The 1-to-1 compatibility analysis can take into account analysis input, including input regarding  
23 the systems 16 to be analyzed, rule set related parameters, workload related parameters,  
24 workload benchmarks and importance factors 88 used to compute overall scores.

1 [00116] The compatibility analysis evaluates the compatibility of every specified system as  
2 source-target pairs on a 1-to-1 basis. This analysis produces a compatibility score for each  
3 system pair so that analyzing a collection of ten (10) systems 16 produces 10x10 scores. The  
4 compatibility analysis is based on the specified rule sets and workload types.

5 [00117] An analysis may be based upon zero or more rule sets and zero or more workload  
6 types, such that at least one rule set or workload type is selected. Example rule sets 28 and  
7 corresponding descriptions are shown in Figure 12, and example workload types 30 and  
8 corresponding descriptions are shown in Figure 13.

9 [00118] The selection of rule sets 28 and workload types 30 for an analysis depends on the  
10 systems 28 and the consolidation strategy to analyze. For example, to assess the consolidation of  
11 a set of UNIX™ systems 16, an analysis may employ the UNIX™ application stacking, location,  
12 maintenance window and ownership rule sets 28, and CPU, memory, disk space, disk I/O and  
13 network I/O workload types 30.

#### 14 **1-to-1 Compatibility Analysis Process Flow**

15 [00119] A process flow diagram of the 1-to-1 compatibility analysis is shown in Figure 14.  
16 The analysis generally comprises four stages.

17 [00120] In the first stage, data referenced by the selected rule sets 28 and workload types 30  
18 for the specified date range are retrieved from the data repository 54 and caches 56, 58 for each  
19 system 16 to be analyzed. This analysis data is saved as a snapshot and can be used for  
20 subsequent analyses.

21 [00121] In the second stage, technical and business related compatibility may be analyzed the  
22 using the specified rule sets 28 and weights. Next, workload compatibility is evaluated based the  
23 specified workload types 30 and input parameters. Finally, the overall compatibility scores are  
24 computed for each pair of systems 16.

25 [00122] Upon completion of the compatibility analysis, the results 20 are provided to the user.  
26 The results 20 include rule item and workload data snapshots, 1-to-1 compatibility score maps



for each rule set 28 and workload type 30 as well as an overall score map. Analysis details for each map may also be provided.

[00123] As noted above, the differential rule sets 28 are used to evaluate the compatibility of systems as they relate to technical and business related constraints. The rule set 28 defines which settings are important for determining compatibility. The rule set 28 typically defines a set of rules which can be revised as necessary based on the specific environment 12. The rule set 28 is thus preferably compiled according to the systems 16 being analysed and prior knowledge of what makes a system 16 compatible with another system 16 for a particular purpose. As will be discussed below, the rule sets 28 are a form of metadata 62.

## Differential Rule Sets

[00124] Further detail regarding the differential rules and differential rule sets 28 is now described making reference to Figures 15 and 16, as also described in co-pending U.S. Patent Application No. 11/535,308 filed on September 26, 2006, and entitled "Method for Evaluating Computer Systems", the contents of which are incorporated herein by reference.

[00125] With respect to the following description of the rule sets 28 and the general application of the rule sets 28 for detecting system incompatibilities by evaluating differences between data parameters of systems 16, the following alternative nomenclature may be used. A target system refers to a system being evaluated, and a baseline system is a system to which the target system is being compared. The baseline and target systems may be the same system 16 at different instances in time (baseline = prior, target = now) or may be different systems 16 being compared to each other. As such, a single system 16 can be evaluated against itself to indicate changes with respect to a datum as well as how it compares to its peers. It will be appreciated that the terms "source system" and "baseline system" are herein generally synonymous, whereby a source system is a type of baseline system.

[00126] Figure 1 illustrates the relationships between system data 18 and the analysis program 10. Data 18 is obtained from the source and target computer systems 16 and is used to analyze the compatibility between the systems 16. In this example, the parameters are evaluated

1 to determine system compatibilities for a consolidation strategy. A distinct data set 18 is  
2 preferably obtained for each system 16 (or instance in time for the same system 16 as required).

3 **[00127]** Rule sets 28 are computer readable and storable so that they may be accessed by the  
4 program 10 and modified if necessary, for use in evaluating the computer systems 16.

5 **[00128]** Rule sets 28 are groupings of rules that represent higher-level considerations such as  
6 business objectives or administrative concerns that are taken into account when reporting on or  
7 analysing the systems 16. In Figure 15, six rules 43, A, B C, D, E and F are grouped into three  
8 rule sets 28, Rule Set 1, 2 and 3. It will be appreciated that there may be any number of rules in  
9 any number of rule sets 28 and those shown in Figure 15 are for illustrative purposes only.

10 **[00129]** Rules evaluate data parameters according to rule definitions to determine  
11 incompatibilities due to differences (or contentious similarities) between the baseline and target  
12 systems. The rule definitions include penalty weights that indicate the importance of the  
13 incompatibility as they relate to the operation of the systems 16. The penalty weights are applied  
14 during an evaluation if the incompatibility is detected. The evaluation may include the  
15 computation of a score or generation of other information indicative of nature of the  
16 incompatibilities between the baseline and target systems.

17 **[00130]** Rules comprised by a rule set 28 may reference common parameters but perform  
18 different tests to identify different forms of incompatibilities that may have different levels of  
19 importance. For example a version four operating system versus a version three operating  
20 system may be considered less costly to remedy and thus less detrimental than a version five  
21 operating system compared to a version one operating system. As can be seen, even though the  
22 operating systems are different in both cases, the nature of the difference can also be considered  
23 and different weights and/or remedies applied accordingly.

24 **[00131]** Rules can also test for similarities that indicate contentions which can result in  
25 incompatibilities between systems. For example, rules can check for name conflicts with respect  
26 to system names, database instance names, user names, etc.

[00132] The flow of data for applying exemplary rule sets 28 is shown in Figure 15. In this example, the system data gathered from a pair of systems 16 are evaluated using three rule sets. The rule engine 90 (see also Figure 7) evaluates the data parameters of the systems 16 by applying rule sets 1, 2 and 3 which comprise of the exemplary rules A, B, C, D, E and F. The evaluation of the rules results in compatibility scores and zero or more matched rule items for each rule set 28. These results can be used for subsequent analyses, such as combining with workload compatibility results to obtain overall compatibility scores.

## Rule Set Specification

[00133] Each rule set 28 has a unique rule set identifier (UUID), rule set name, rule set description, rule set version, rule set type (e.g. controlled by system 10 or by user), and a rule set category (e.g. generic rule set categorization such as business, configuration etc.).

[00134] As described above, each rule set 28 is also comprised of one or more rules.

## Rule Definition

[00135] A rule is conceptually a form of metadata 62 that specifies a parameter, tests to detect an incompatibility between the baseline and target systems 16, the relative importance of the incompatibility, and the costs associated with the remediating the incompatibility. Each rule is defined by a fixed number of fields as listed in Table 1 below.

Field	Description
Name	Rule name
Description	Rule description
Data Query Type	Query type to get data parameter (e.g. URI, Attribute, Alias)
Query Value	Data query specification based on query type
Baseline Test	Baseline test specification
Target Test	Target test specification
Weight	Rule penalty weight
Mutex Flag	Y/N – This flag is used at multiple levels as described below
Match Flag	Rule match name referenced by suppress flag
Suppress Flag	Rule dependency expression to determine whether to suppress rule

Field	Description
Remediation Costs	Estimated remediation costs for rule item if true.
Enabled Flag	True/False

*Table 1 – Rule Item Field Specification*

**[00136]** The name and description fields are lay descriptions of the condition or discrepancy detected by the rule. These fields are used to provide management-level summaries when processing rule sets 28. The fields can provide as much or as little information as required by the application.

## **Rule Query Specification**

**[00137]** The data query type and value identify the data parameter to be evaluated by the rule. The query type specifies whether the rule applies to audited data directly (UriQuery), normalized values (AliasQuery) or attributes (AttrQuery). The query value specifies the parameter specification based on the query type. For UriQuery, AliasQuery and AttrQuery types, query values specify URI, alias and attribute names, respectively.

**[00138]** The URI value conforms to a URI-like format used for accessing parameters from the native audit data model. The URI can specify the module, object and property of the data object or property that is being requested. The optional URI fragment (i.e. the portion after the “#” symbol) specifies the specific object instance (table row) that is being evaluated, with “\*” denoting a wildcard that matches all instances.

## **Rule Match Test**

**[00139]** If specified, the baseline field represents the literal value that would need to match the value of the object/property on the source system in order for the rule to match. For objects and object instances, the keywords “absent” and “present” are preferably used to match cases where that object is absent or present respectively. Similar to the baseline field, the target field allows a literal match against the value of the object/property on the target system. The target field also supports the absent/present specifiers. For numeric properties, relational operators (>, <, =, >=, <=, <>)

<, =, !=) can be used to cause the rule to trigger if the target value has the specified relationship with the source value.

[00140] In order to apply a rule to a target/baseline pair, the following test specification can be followed as shown in Table 2.

	Target	Baseline	Description
1			Values are different
2	!=		Values are different
3		!=	Values are different
4	>	ANY	Target > Baseline
5	<	ANY	Target < Baseline
6	=		Values are the same
7		=	Values are the same
8	~		Values are similar
9		~	Values are similar
10	X	Y	Target = X and Baseline = Y
11	X		Target = X and Baseline != X
12		Y	Target != Y and Baseline = Y

Table 2 – Target/Baseline Test Specification

[00141] The rule test specifications can be extended to include such things as regular expression pattern matching, logical test expressions (using AND, OR), etc. as described below.

#### Rule Weight

[00142] The weight field specifies the relative importance of that property and combination of source/target values (if specified) in regard to the overall context of the comparison. Higher values indicate that the condition detected by the rule has a high impact on the target environment, with 100% being an “absolute constraint” that indicates complete incompatibility.

#### Mutex Flag

[00143] The mutex flag field can be used to avoid multiple penalties that would otherwise skew the scores. A “Y” in the mutex flag field specifies that multiple matches of the same rule

43 will incur only a single penalty on the overall score (as specified in the weight field), as opposed to multiple accumulating penalties (which is the default behaviour).

[00144] The mutex flag can be interpreted at multiple levels. When comparing a target and source system, should the rule specifier expand to a list (e.g. software list), the rule is evaluated for each instance. In this case, if the flag is a “Y”, the score is penalized by the rule weight a maximum of one time, even if the rule was true for multiple instances. If the flag is “N”, the score should be penalized for each instance that was evaluated to be true. Furthermore, when computing multi-dimensional compatibility scores (multiple sources transferred to a single target), the calculation is based on the union of the rule items that were true. In this case, the flag is used to determine whether to penalize the multi-stack score once or for each unique rule instance. The multi-dimensional compatibility analysis is described in detail below.

## **Rule Dependency**

[00145] The match flag field enables an optional symbolic flag to be “set” when a rule matches, and which can subsequently be used to suppress other rules (through the “Suppress Flags” field). This effectively allows rule dependencies to be modeled in the rule set 28. The suppress flag field allows symbolic flags (as specified in the “Match Flag” field) to be used to suppress the processing of rules. This allows specific checks to be skipped if certain higher-level conditions exist. For example, if the operating systems are different, there is no need to check the patches. It should be noted that this allows any arbitrary logic to be constructed, such as how logic can be built from NAND gates.

[00146] The remediation cost field is preferably optional. The remediation field represents the cost of “fixing” the system(s) (i.e. eliminating the condition or discrepancy detected by the rule 43). When analyzing differences between (or changes to) IT systems this is used to represent hardware/software upgrade costs, administrative costs and other costs associated with making the required changes to the target systems. The calculations behind this field vary based on the nature of the system and the parameter that would need to be added, upgraded etc.

1 [00147] Each rule can include a true/false enable flag for enabling and disabling the rule  
2 item.

### 3 Rule Set Example

4 [00148] Figure 16 provides an example rule set 28, which includes a number of rules. The  
5 following refers to the number indicated in the leftmost column of Figure 16.

6 [00149] Rule 1 scrutinizes the normalized (AliasQuery) representation of the operating  
7 systems (e.g. Windows™, Solaris™, AIX™, Linux™, etc.) on both the source and target  
8 systems and heavily penalizes cases where these are different as evident from the high weight  
9 factor (70%). Rule 2 penalizes systems that have different operating system versions (e.g.  
10 Windows™ NT vs Windows™ 2000), and is suppressed (i.e. not processed) in cases where the  
11 systems have different overall operating systems (as detected in the previous rule). Rule 3  
12 detects if systems are in different time zones. Rule 4 penalizes combinations of systems where  
13 the target has less memory than the source (this is what is referred to as a directional rule, which  
14 can give differing results if sources and targets are reversed, e.g. asymmetric results). Rule 5  
15 operates directly against audit data and detects cases where the operating system patch level  
16 differs. This rule is not processed if either the operating system or the operating system version  
17 are different (since this renders the comparison of patches meaningless).

18 [00150] Rule 6 scrutinizes the lists of all patches applied to the source and target systems and  
19 penalizes cases where they differ. The mutex flag is set, indicating that the penalty is applied  
20 only once, no matter how many patch differences exist. This rule is ignored in cases where  
21 either the operating system or operating system version are different. Rule 7 penalizes system  
22 combinations of servers that are running the same OS but are configured to run a different  
23 number of kernel bits (e.g. 64-bit vs 32-bit). Rule 8 penalizes combinations where there are  
24 kernel parameters defined on the source that are not defined on the target. This rule is not  
25 applied if the operating systems are different.

26 [00151] Rule 9 scrutinizes a specific kernel setting (SHMMAX, the setting that specifies how  
27 much shared memory a system can have) and penalizes combinations where it is set to a lower

1 value on the target than it is on the source system. Rule 10 penalizes combinations of systems  
2 that are running different database version, e.g. Oracle™ 9 vs. Oracle™ 8. Rule 11 penalizes  
3 combinations of systems that are running different versions of Oracle™. Rule 11 is suppressed  
4 if the more specific Rule 10 is true. It should be noted that the remediation cost is relatively  
5 high, owing to the fact that it will take a software upgrade to eliminate this discrepancy. In some  
6 cases the remediation cost can be low where the upgrade is less expensive. Rule 12 penalizes  
7 combinations of systems that are running different versions of Apache. It should be noted that  
8 the remediation cost is relatively low, as apache is an open source product and the cost of  
9 upgrade is based on the hourly cost of a system administrator and how long it will take to  
10 perform the upgrade.

11 **[00152]** Rule 13 scrutinizes a windows-specific area of the audit data to determine if the  
12 source and target systems are running different service pack levels. It should be noted that this  
13 rule closely mirrors rule 5, which uses a rule specifier that scrutinizes the UNIX™/Linux™ area  
14 of the audit data. Rule 14 scrutinizes the lists of all hotfixes applied to the source and target  
15 systems and penalizes cases where they differ. This rule closely mirrors rule 6, which scrutinizes  
16 patches on UNIX™ and Linux™. Rule 15 detects differing startup commands between systems.  
17 Rule 16 is a rule to detect differing Paths between systems, and rule 17 detects differing System  
18 Paths between systems.

19 **[00153]** Rule 18 penalizes system combinations where there are services installed on the  
20 source that are not installed on the target. This rule has the mutex flag set, and will therefore  
21 only penalize a system combination once, no matter how many services are missing. Rule 19  
22 penalizes system combinations where there are services started on the source that are not started  
23 on the target. It should be noted that both the weight and the remediation cost are lower than the  
24 previous rule, owing to the fact that it is generally easier and less expensive to start a service than  
25 install it. Finally, rule 20 penalizes combinations where the target system is missing the virus  
26 scanner software.

27 **[00154]** It will be appreciated that the above described rules and rule set 28 are shown for  
28 illustrative purposes only and that any combination of rules can be used to achieve specific goals.



For example, rules that are applicable to the OS can be grouped together to evaluate how a system compares to its peers. Similarly, rules pertaining to database, Java applications etc. can also be grouped.

[00155] As discussed above, Figure 12 provides a table listing several additional example rule sets and associated descriptions.

[00156] The system consolidation analysis computes the compatibility of a set of systems based not only on technical and workload constraints as exemplified above, but also business constraints. The business constraints can be expressed in rule sets 28, similar to the technical constraints discussed above.

## Rule Specification Enhancements

[00157] In another embodiment, the rule test specification shown above in Table 1 can be extended to support the "NOT" expression, e.g. "!X" where "!" indicates not equal to; and to support regular expression patterns, e.g. "rx:Pattern" or "!rx:Pattern". An example is shown in Table 3 below.

Target	Baseline	Description
X	!Y	Target = X and Baseline != Y
!X		Target != X
	!Y	Baseline != Y
rx:P1	!rx:P2	Target matches P1 and Baseline does not match P2
rx:P1		Target matches P1
	!rx:P2	Baseline does not match P2
X	rx:P2	Target = X and Baseline matches P2
!X	rx:P2	Target != X and Baseline matches P2

*Table 3: Extended Rule Test Specification*

[00158] In yet another embodiment, an advanced rule set specification may also be used to support more flexible rule logic by considering the following rule item specification shown in Table 4.

Field	Description
Name	As before

21635070.1

Field	Description
Description	As before
Data Query Type	As before
Query Value	As before
<b>Test</b>	Test expression
Weight	As before
Mutex Flag	As before
Match Flag	As before
<b>Match Test</b>	Rule dependency expression to determine whether to perform the corresponding match action
<b>Match Action</b>	Action (run/suppress/stop) to execute based on match test result
Remediation Costs	As before
Enabled Flag	As before

*Table 4: Advanced Rule Item Specification*

[00159] The **highlighted** fields above are incorporated to include a test expression and match test and corresponding match action to be executed according to the match test result.

[00160] In the advanced rule specification, the following elements are supported: variables, such as target, source etc.; regular expressions, e.g. regex("pattern"); constants, such as quoted string values "lala", "1" etc.; test operators, =, !=, <, <=, >, >=, ~, !~; logical operators, AND &&, OR, ||; and logical operator precedence, ( ) for nested expressions.

[00161] For example:

[00162] Target != Source;

[00163] Target >= Source && > "1024";

[00164] Target = regex("Windows") && Source != regex("Solaris|AIX"); and

[00165] (Target = "X" && Source = "Y") || (Target = "A" && Source = "B").

[00166] The test expression supports the following elements: match flag names; logical operators, AND &&, OR, ||; test operator, NOT !; and logical operator precedence, ( ) for nested expressions. Examples assuming match flags of A, B, C, D, etc. are as follows:

1 [00167] A

2 [00168] A || B

3 [00169] A && B

4 [00170] (A && !B) || C

5 [00171] The match action specifies the action to perform if the match test rule is true.

6 Possible actions are: run, evaluate the rule item; suppress, do not evaluate the rule item; and stop.

7 [00172] Where basic and advanced rule sets are available for the same analysis program, there  
8 are a number of options for providing compatibility. The rule set specification can be extended  
9 to include a property indicating the minimum required rule engine version that is compatible  
10 with the rule set. In addition, the basic rule sets can be automatically migrated to the advanced  
11 rule set format since the advanced specification provides a super set of functionality relative to  
12 the basic rule set specification. It will be appreciated that as new rules and rule formats are  
13 added, compatibility can be achieved in other ways so long as legacy issues are considered where  
14 older rule versions are important to the analysis.

## 15 **1-to-1 Rule-based Compatibility Analysis**

16 [00173] An exemplary process flow for a rule-based compatibility analysis is shown in greater  
17 detail in Figures 17 and 18. When analyzing system compatibility, the list of target and source  
18 systems 16 are the same. The compatibility is evaluated in two directions, e.g. for a Server A  
19 and a Server B, migrating A to B is considered as well as migrating B to A.

20 [00174] Turning first to Figure 17, for each rule set R (R = 1 to M where M is the number of  
21 rule sets) and for each target system T (T = 1 to N where N is the number of systems), the rule  
22 engine 90 first looks at each source system S (S = 1 to N). If the source=target then the  
23 configuration compatibility score for that source is set to 100, no further analysis is required and  
24 the next pair is analyzed. If the source and target are different, the rules are evaluated against the  
25 source/target pair to compute the compatibility score, remediation cost and to compile the

1 associated rule details. Estimated remediation costs are optionally specified with each rule item.  
2 As part of the rule evaluation and subsequent compatibility score calculation, if a rule is true, the  
3 corresponding cost to address the deficiency is added to the remediation cost for the pair of  
4 systems 16 being analysed.

5 **[00175]** The evaluation of the rules is shown in Figure 18. The evaluation of the rules  
6 considers the snapshot data 18 for the source system and the target system, as well as the  
7 differential rule set 28 that being applied. For each rule in the set 28, the data referenced by the  
8 rule is obtained for both the target and source. The rule is evaluated by having the rule engine 90  
9 compare the data. If the rule is not true (i.e. if the systems 16 are the compatible according to the  
10 rule definition) then the data 18 is not considered in the compatibility score and the next rule is  
11 evaluated. If the rule is true, the rule details are added to an intermediate result. The  
12 intermediate result includes all true rules.

13 **[00176]** Preferably, a suppression tag is included with each rule. As discussed above, the  
14 suppression tag indicates other rules that are not relevant if that rule is true. The suppression flag  
15 allows the program 10 to avoid unnecessary computations. A mutex flag is also preferably used  
16 to avoid unfairly reducing the score for each true rule when the rules are closely affected by each  
17 other.

18 **[00177]** Once each rule has been evaluated, a list of matched rules is created by removing  
19 suppressed rule entries from the intermediate results based on rule dependencies, which are  
20 defined by rule matching and suppression settings (e.g. match flags and suppression tags). The  
21 compatibility score for that particular source/target pair is then computed based on the matched  
22 rules, weights and mutex settings. Remediation costs are also calculated based on the cost of  
23 updating/upgrading etc. and the mutex settings.

24 **[00178]** Turning back to Figure 17, the current target is then evaluated against all remaining  
25 sources and then the next target is evaluated. As a result, an N x N map 32 can be created that  
26 shows a compatibility score for each system against each other system. The map 32 can be  
27 sorted by grouping the most compatible systems. The sorted map 32 is comprised of every

source/target combination and thus provides an organized view of the compatibilities of the systems 16.

[00179] Preferably, configuration compatibility results are then generated for each rule set 28, comprising the map 32 (e.g. Figure 4) and for each source-target pair details available pertaining to the configuration compatibility scoring weights, remediation costs and applicable rules. The details can preferably be pulled for each source/target pair by selecting the appropriate cell 92 (see Figure 19).

### 1-to-1 Workload Compatibility Analysis

[00180] The workload compatibility analysis evaluates the compatibility of each source-target pair with respect to one or more workload data types 30. The analysis employs a workload stacking model to combine the source workloads onto the target system. The combined workloads are then evaluated using threshold and a scoring algorithm to calculate a compatibility score for each workload type.

### Workload Data Types and Benchmarks

[00181] System workload constraints must be assessed when considering consolidation to avoid performance bottlenecks. Workload types representing particularly important system resources include % CPU utilization, memory usage, disk space used, disk I/O throughput and network I/O throughput. The types of workload analyzed can be extended to support additional performance metrics. As noted above, example workload types are summarized in the table shown in Figure 13.

[00182] Each workload type can be defined by the properties listed in Table 5 below:

Property	Description	Examples
Name	Workload key name	CPU_Utilization
Display Name	Workload display name for UI	CPU Utilization
Benchmark	Benchmark type corresponding to workload	cpu

Property	Description	Examples
Type	type	
Alias Name	Alias to get workload values from repository	CpuDays
Alias File	Alias file containing above alias	cpu_workload_alias.xml
Description	Short description of workload type	CPU Utilization
Unit	Unit of workload value	%
Test as percent?	Boolean flag indicating whether to test the workload against a threshold as a percentage (true) or as an absolute value (false)	true

*Table 5: Workload Type Definition*

[00183] Workload values can be represented as percentages (e.g. %CPU used) or absolute values (e.g. disk space used in MB, disk I/O in MB/sec).

[00184] The term workload benchmark refers to a measure of a system's capability that may correspond to one or more workload types. Workload benchmarks can be based on industry benchmarks (e.g. CINT2000 for processing power) or the maximum value of a system resource (e.g. total disk space, physical memory, network I/O bandwidth, maximum disk I/O rate). Benchmarks can be used to normalize workload types that are expressed as a percentage (e.g. %CPU used) to allow direct comparison of workloads between different systems 16.

[00185] Benchmarks can also be used to convert workload types 30 that are expressed as absolute values (e.g. disk space used in MB) to a percentage (e.g. % disk space used) for comparison against a threshold expressed as a percentage. Each benchmark type can be defined by the following in Table 6:

Property	Description	Example
Name	Benchmark name	cpu
Default value	Default benchmark value if not resolved by the other means (optional)	<none>
Alias name	Alias to get benchmark value for specific	cpuBenchmark

Property	Description	Example
	system (optional)	
Alias file	File containing alias specified above	benchmark_alias.xml
Attribute name	System attribute value to lookup to get benchmark value (optional)	CINT2000
Use alias first	Boolean flag indicating whether to try the alias or the attribute lookup first	true

Table 6: Workload Benchmark Definition

[00186] System benchmarks can normalize workloads as follows. For systems X and Y, with CPU benchmarks of 200 and 400 respectively (i.e. Y is 2x more powerful than X), if systems X and Y have average CPU utilizations of 10% and 15% respectively, the workloads can be normalized through the benchmarks as follows. To normalize X's workload to Y, multiply X's workload by the benchmark ratio X/Y, i.e.  $10\% \times 200/400 = 5\%$ .

[00187] Stacking X onto Y would then yield a total workload of  $5\% + 15\% = 20\%$ . Conversely, stacking Y onto X would yield the following total workload:  $10\% + 15\% \times 400/200 = 40\%$ .

## Workload Data Model

[00188] As discussed above, workload data is collected for each system through various mechanisms including agents, standard instrumentation (e.g. Windows Performance Monitor™, UNIX™ System Activity Reporter), custom scripts, third party performance monitoring tools, etc. Workload data is typically collected as discrete time series data. Higher sample frequencies provide better accuracy for the analysis (5 minute interval is typical). The workload data values should represent the average values over the sample period rather than instantaneous values. An hour of CPU workload data for three fictional systems (A, B and C) is listed below in Table 7:

Timestamp	%CPU used (A)	%CPU used (B)	%CPU used (C)
03/01/07 00:00:00	10	0	0

Timestamp	%CPU used (A)	%CPU used (B)	%CPU used (C)
03/01/07 00:05:00	12	0	0
03/01/07 00:10:00	18	10	4
03/01/07 00:15:00	22	10	6
03/01/07 00:20:00	25	15	7
03/01/07 00:25:00	30	25	8
03/01/07 00:30:00	30	35	12
03/01/07 00:35:00	35	39	15
03/01/07 00:40:00	39	45	19
03/01/07 00:45:00	41	55	21
03/01/07 00:50:00	55	66	28
03/01/07 00:55:00	80	70	30

*Table 7: Sample Workload Data (Time series)*

[00189] Data from different sources may need to be normalized to common workload data types to ensure consistency with respect to what and how the data is measured. For example, %CPU usage may be reported as Total %CPU utilization, %CPU idle, %CPU system, %CPU user, %CPU I/O, etc. Disk utilization may be expressed in different units such as KB, MB, blocks, etc.

[00190] The time series workload data can be summarized into hourly quartiles. Specifically, the minimum, 1<sup>st</sup> quartile, median, 3<sup>rd</sup> quartile, maximum, and average values are computed for each hour. Based on the workload data from the previous example, the corresponding summarized workload statistics are listed below in Table 8:

System	Hour	%CPU (Avg)	%CPU (Min)	%CPU (Q1)	%CPU (Q2)	%CPU (Q3)	%CPU (Max)
A	0	33.1	10	20	30	40	80
B	0	30.8	0	10	30	50	70



C	0	12.5	0	5	10	20	30
---	---	------	---	---	----	----	----

*Table 8: Summarized Workload Data (Quartiles) for Hour 0*

**[00191]** The compatibility analysis for workload uses the hourly quartiles. These statistics allow the analysis to emphasize the primary operating range (e.g. 3<sup>rd</sup> quartile) while reducing sensitivity to outlier values.

### **Workload Data Extraction**

**[00192]** Workload data is typically collected and stored in the workload data cache 58 for each system 16 for multiple days. At least one full day of workload data should be available for the analysis. When analyzing workloads, users can specify a date range to filter the workload data under consideration. A representative day is selected from this subset of workload data for the analysis. The criteria for selecting a representative day should be flexible. A preferable default assessment of the workload can select the worst day as the representative day based on average utilization. A less conservative assessment may consider the N<sup>th</sup> percentile (e.g. 95<sup>th</sup>) day to eliminate outliers. Preferably, the worst days (based on daily average) for each system and for each workload type are chosen as the representative days.

**[00193]** The data extraction process flow for the workload compatibility analysis is shown in Figure 20. Preferably, the workload data cache 58 includes data obtained during one or more days. For each system 16 in the workload data set, for each workload data type 30, get the workload data for the specified date range, determine the most representative day of data, (e.g. if it is the worst day) and save it in the workload data snapshot. In the result, a snapshot of a representative day of workload data is produced for each system 16.

### **Workload Stacking**

**[00194]** To evaluate the compatibility of one or more systems with respect to server consolidation, the workloads of the source systems are combined onto the target system. Some types of workload data are normalized for the target system. For example, the %CPU utilization is normalized using the ratio of target and source CPU processing power benchmarks. The

21635070.1

consolidated workload for a specific hour in the representative day is approximated by combining the hourly quartile workloads.

[00195] There are two strategies for combining the workload quartiles, namely original and cascade. The original strategy simply adds like statistical values (i.e. maximum, third quartile, medians, etc.) of the source systems to the corresponding values of the target system. For example, combining the normalized CPU workloads of systems A and B onto C (from Table 8), the resulting consolidated workload statistics for hour 0 are:

$$\text{[00196] Maximum} = \text{Max}_A + \text{Max}_B + \text{Max}_C = 180\%$$

$$\text{[00197] 3}^{\text{rd}} \text{ Quartile} = Q3_A + Q3_B + Q3_C = 110\%$$

$$\text{[00198] Median} = Q2_A + Q2_B + Q2_C = 70\%$$

$$\text{[00199] 1}^{\text{st}} \text{ Quartile} = Q1_A + Q1_B + Q1_C = 35\%$$

$$\text{[00200] Minimum} = \text{Min}_A + \text{Min}_B + \text{Min}_C = 10\%$$

[00201] The resulting sums can exceed theoretical maximum capacity/benchmark values (e.g. >100% CPU used).

[00202] The cascade strategy processes the statistical values in descending order, starting with the highest statistical value (i.e. maximum value). The strategy adds like statistical values as with original, but may clip the resulting sums if they exceed a configurable limit and cascades a portion of the excess value to the *next statistic* (i.e. the excess of sum of the maximum values is cascaded to 3<sup>rd</sup> quartile). The relevant configuration properties for the cascade calculation are listed below in Table 9.

Property	Description	Example
wci.cascade_overflow	Boolean flag indicating whether to apply cascade (true) or original (false) strategy	true/false
wci.cascade_overflow_proportion	Fraction of overflow to cascade to next value	0.5

Property	Description	Example
wci.clip_type	Flag indicating whether to use benchmark value (max) or analysis workload threshold (limit) as the clipping limit	limit max
wci.clip_limit_ratio	If clip_type is limit, this is the ratio to apply to workload limit to determine actual clipping limit	2

Table 9: Workload Cascade Configuration

The following example applies cascade calculation to the example workload data from Table 2 with the following settings:

wci.cascade\_overflow = true

wci.cascade\_overflow\_proportion = 0.5

wci.clip\_type = max

The *max* clip type indicates that the clipping level is the maximum value of the workload type. In this example, the clipping level is 100% since the workload type is expressed as a percentage (%CPU). The overflow proportion indicates that 0.5 of the excess amount above the clipping level should be cascaded to the next statistic. The consolidated workload statistics are computed as follows:

Max<sub>Original</sub> = Max<sub>A</sub> + Max<sub>B</sub> + Max<sub>C</sub>

= 180%

Max<sub>Clipped</sub> = Minimum of (Max<sub>Original</sub>, clipping level)

= Minimum of (180, 100)

= 100%

Max<sub>Excess</sub> = Max<sub>Original</sub> - Max<sub>Clipped</sub>

1	[00214]	=	180% - 100%
2	[00215]	=	80%
3	[00216] MaxOverflow	=	MaxExcess * wci.cascade_overflow_proportion
4	[00217]	=	80% * 0.5
5	[00218]	=	40%
6	[00219] Q3Original	=	Q3 <sub>A</sub> + Q3 <sub>B</sub> + Q3 <sub>C</sub>
7	[00220]	=	110%
8	[00221] Q3Cascade	=	Q3Original + MaxOverflow
9	[00222]	=	110% + 40%
10	[00223]	=	150%
11	[00224] Q3Clipped	=	Minimum of (Q3Cascade, clipping level)
12	[00225]	=	Minimum of (150, 100)
13	[00226]	=	100%
14	[00227] Q3Excess	=	50%
15	[00228] Q3Overflow	=	25%
16	[00229] Q2Original	=	70%
17	[00230] Q2Cascade	=	Q2Original + Q3Overflow
18	[00231]	=	70% + 25%
19	[00232]	=	95%

- 1    [00233]     $Q2_{Clipped}$                       =        Minimum of ( $Q2_{Cascade}$ , clip level)
- 2    [00234]    =        Minimum of (95, 100)
- 3    [00235]    =        95%
- 4    [00236]     $Q2_{Overflow}$                       =        0%
- 5    [00237]     $Q1_{Original}$                       =        35%
- 6    [00238]     $Q1_{Cascade}$                       =        35%
- 7    [00239]     $Q1_{Clipped}$                       =        35%
- 8    [00240]     $Q1_{Overflow}$                       =        0%
- 9    [00241]     $Min_{Original}$                       =        10%
- 10   [00242]     $Min_{Clipped}$                       =        10%
- 11   [00243]    The consolidated statistics for the above example are summarized in the following
- 12   Table 10. The *clipped* values are net results of the analysis, namely the new answer.

Statistic	Original	Cascade	<i>Clipped</i>	Excess	Overflow
Max	180	180	<b>100</b>	80	40
Q3	110	150	<b>100</b>	50	25
Q2	70	95	<b>95</b>	0	0
Q1	35	35	<b>35</b>	0	0
Min	10	10	<b>10</b>	0	0

13    *Table 10: Cascaded Statistics (Example 1)*

- 14   [00244]    Similarly, the following example applies cascade calculation to the example workload
- 15   data from Table 7 with the following settings:

- 16   [00245]    `wci.cascade_overflow`                      =        `true`

21635070.1

1 [00246] wci.cascade\_overflow\_proportion = 0.5

2 [00247] wci.clip\_type = limit

3 [00248] wci.clip\_limit\_ratio = 2

4 [00249] This example specifies a *limit* clip type to indicate that the clipping level is based on  
5 the analysis threshold for the workload type. The *clip\_limit\_ratio* specifies the ratio to apply to  
6 the threshold to calculate the actual clipping level. For instance, if the threshold is 80% and the  
7 clip limit ratio is 2, the clipping level is 160%. The consolidated workload statistics based on the  
8 above settings listed below in Table 11.

Statistic	Original	Cascade	<i>Clipped</i>	Excess	Overflow
Max	180	180	<i>160</i>	20	10
Q3	110	120	<i>120</i>	0	0
Q2	70	70	<i>70</i>	0	0
Q1	35	35	<i>35</i>	0	0
Min	10	10	<i>10</i>	0	0

9 *Table 11: Cascaded Statistics (Example 2)*

## 10 **Workload Compatibility Scoring**

11 [00250] Workload compatibility scores quantify the compatibility of consolidating one or  
12 more source systems onto a target system. The scores range from 0 to 100 with higher scores  
13 indicating better compatibility. The scores are computed separately for each workload type 30  
14 and are combined with the system configuration and business-related compatibility scores to  
15 determine the overall compatibility scores for the systems 16. The workload scores are based on  
16 the following: combined system workload statistics at like times and worst case, user-defined  
17 workload thresholds, penalty calculation, score weighting factors, and workload scoring formula.

18 [00251] Workloads are assessed separately for two scenarios: like-times and worst case. The  
19 like times scenario combines the workload of the systems at like times (i.e. same hours) for the

21635070.1

representative day. This assumes that the workload patterns of the analyzed systems are constant. The worst case scenario time shifts the workloads for one or more systems to determine the peak workloads. This simulates the case where the workload patterns of the analyzed systems may occur earlier or be delayed independently. The combined workload statistics (maximum, 3<sup>rd</sup> quartile, median, 1<sup>st</sup> quartile and minimum) are computed separately for each scenario.

[00252] For a specific analysis, workload thresholds are specified for each workload type. The workload scores are penalized as a function of the amount the combined workload exceeds the threshold. Through the workload type definition (Table 6), the workload data (Table 7) and corresponding thresholds can be specified independently as percentages or absolute values. The workload data type 30 is specified through the *unit* property and the threshold data type is specified by the *test as percent* flag. The common workload/threshold data type permutations are handled as follows.

[00253] If the workload is expressed as a percentage and *test as percent* is true (e.g. %CPU), normalize workload percentage using the benchmark and compare as percentages.

[00254] If the workload is expressed as an absolute value and *test as percent* is true (e.g. disk space), convert the workload to a percentage using benchmark and compare as percentages.

[00255] If workload unit is expressed as an absolute value and *test as percent* if false (e.g. network I/O), compare workload value against threshold as absolute values.

[00256] A penalty value ranging from 0 to 1 can be calculated for each workload statistic and for each scenario as a function of the threshold and the clipping level. The penalty value is computed as follows:

[00257] If Workload  $\leq$  Threshold,

[00258] Penalty = 0

[00259] If Workload  $\geq$  Clipping Level,

1 [00260] Penalty = 1

2 [00261] If Threshold < Workload < Clipping Level,

3 [00262] Penalty = (Workload Value - Threshold) / (Clipping level – Threshold)

4 [00263] Using Example 2 from above (threshold = 80%, clipping level = 160%), the sliding  
5 scale penalty values are computed as follows:

6 [00264]  $\text{Penalty}_{\text{Max}} = (160 - 80) / (160 - 80)$

7 [00265] = 1

8 [00266]  $\text{Penalty}_{\text{Q3}} = (120 - 80) / (160 - 80)$

9 [00267] = 0.5

10 [00268]  $\text{Penalty}_{\text{Q2}} = 0$  [since 70 < 80]

11 [00269]  $\text{Penalty}_{\text{Q1}} = 0$  [since 35 < 80]

12 [00270]  $\text{Penalty}_{\text{Min}} = 0$  [since 10 < 80]

13

14 [00271] The workload score is composed of the weighted penalty values. Penalty weights can  
15 be defined for each statistic and scenario as shown in Table 12 below.

Statistic	Scenario	Property	Example
Maximum	Like Times	wci.score.max_like_times	0.2
Maximum	Worst Times	wci.score.max_worst_times	0.1
3 <sup>rd</sup> Quartile	Like Times	wci.score.q3_like_times	0.4
3 <sup>rd</sup> Quartile	Worst Times	wci.score.q3_worst_times	0.3
Median	Like Times	wci.score.q2_like_times	0
Median	Worst Times	wci.score.q2_worst_times	0
1 <sup>st</sup> Quartile	Like Times	wci.score.q1_like_times	0
1 <sup>st</sup> Quartile	Worst Times	wci.score.q1_worst_times	0
Minimum	Like Times	wci.score.min_like_times	0
Minimum	Worst Times	wci.score.min_worst_times	0

16



Table 12: Score Weighting Factors

[00272] The weights are used to compute the workload score from the penalty values. If the sum of the weights exceeds 1, the weights should be normalized to 1.

[00273] The actual score is computed for a workload type by subtracting the sum of the weighted penalties from 1 and multiplying the result by 100:

$$[00274] \quad \text{Score} = 100 * (1 - \text{Sum}(\text{Weight} * \text{Penalty}))$$

[00275] Using the previous example and assuming that the like times are the same as the worst times, the score is calculated as follows:

$$[00276] \quad \text{Score} = 100 * (1 - (\text{Weight}_{\text{Max Worst}} * \text{Penalty}_{\text{Max Worst}} + \text{Weight}_{\text{Max Like}} * \text{Penalty}_{\text{Max Like}} +$$

$$[00277] \quad \text{Weight}_{\text{Q3 Worst}} * \text{Penalty}_{\text{Q3 Worst}} + \text{Weight}_{\text{Q3 Like}} * \text{Penalty}_{\text{Q3 Like}} +$$

$$[00278] \quad \text{Weight}_{\text{Q2 Worst}} * \text{Penalty}_{\text{Q2 Worst}} + \text{Weight}_{\text{Q2 Like}} * \text{Penalty}_{\text{Q2 Like}} +$$

$$[00279] \quad \text{Weight}_{\text{Q1 Worst}} * \text{Penalty}_{\text{Q1 Worst}} + \text{Weight}_{\text{Q1 Like}} * \text{Penalty}_{\text{Q1 Like}} +$$

$$[00280] \quad \text{Weight}_{\text{Min Worst}} * \text{Penalty}_{\text{Min Worst}} + \text{Weight}_{\text{Min Like}} * \text{Penalty}_{\text{Min Like}}))$$

$$[00281] \quad = 100 * (1 - (0.1*1 + 0.2*1 + 0.3*0.5 + 0.4*0.5))$$

$$[00282] \quad = 30$$

## 1-to-1 Workload Compatibility Analysis Process Flow

[00283] A flow chart illustrating a workload compatibility analysis is shown in Figure 21.

When analyzing 1-to-1 workload compatibility, the list of target and source systems 16 is the same. The compatibility is evaluated in two directions, e.g. for Server A and Server B, migrating A to B is considered as well as migrating B to A.

[00284] The workload analysis considers one or more workload types, e.g. CPU busy, the workload limits 94, e.g. 75% of the CPU being busy, and the system benchmarks 96, e.g. relative CPU power. Each system 16 in the workload data set is considered as a target (T = 1 to N) and

21635070.1

1 compared to each other system 16 in the data set 18 as the source ( $S = 1$  to  $N$ ). The analysis  
2 engine 64 first determines if the source and target are the same. If yes, then the workload  
3 compatibility score is set to 100 and no additional analysis is required for that pair. If the source  
4 and target are different, the system benchmarks are then used to normalize the workloads (if  
5 required). The normalized source workload histogram is then stacked on the normalized target  
6 system.

7 **[00285]** System benchmarks can normalize workloads as follows. For systems X and Y, with  
8 CPU benchmarks of 200 and 400 respectively (i.e. Y is 2x more powerful than X), if systems X  
9 and Y have average CPU utilization of 10% and 15% respectively, the workloads can be  
10 normalized through the benchmarks as follows. To normalize X's workload to Y, multiply X's  
11 workload by the benchmark ratio  $X/Y$ , i.e.  $10\% \times 200/400 = 5\%$ . Stacking X onto Y would then  
12 yield a total workload of  $5\% + 15\% = 20\%$ . Conversely, stacking Y onto X would yield the  
13 following total workload:  $10\% + 15\% \times 400/200 = 40\%$ .

14 **[00286]** Using the stacked workload data, the workload compatibility score is then computed  
15 for each workload type as described above.

16 **[00287]** Each source is evaluated against the target, and each target is evaluated to produce an  
17  $N \times N$  map 34 of scores, which can be sorted to group compatible systems (see Figure 5).  
18 Preferably, a workload compatibility results is generated that includes the map 34 and workload  
19 compatibility scoring details and normalized stacked workload histograms that can be viewed by  
20 selecting the appropriate cell 92 (see Figure 22). The workload compatibility results are then  
21 combined with the rule-based compatibility results to produce the overall compatibility scores,  
22 described below.

### 23 **1-to-1 Overall Compatibility Score Calculation**

24 **[00288]** The results of the rule and workload compatibility analyses are combined to compute  
25 an overall compatibility score for each server pair. These scores preferably range from 0 to 100,  
26 where higher scores indicate greater compatibility and 100 indicating complete or 100%  
27 compatibility.

21635070.1

[00289] As noted above, the analysis input can include importance factors. For each rule set 28 and workload type 30 included in the analysis, an importance factor 88 can be specified to adjust the relative contribution of the corresponding score to the overall score. The importance factor 88 is an integer, preferably ranging from 0 to 10. A value of 5 has a neutral effect on the contribution of the component score to the overall score. A value greater than 5 increase the importance whereas a value less than 5 decreases the contribution.

[00290] The overall compatibility score for the system pair is computed by combining the individual compatibility scores using a formula specified by an overlay algorithm which performs a mathematical operation such as multiply or average, and the score is recorded.

[00291] Given the individual rule and workload compatibility scores, the overall compatibility score can be calculated by using the importance factors as follows for a "multiply" overlay:

$$O = 100 * \frac{100 - (100 - S_1) * \frac{F_1}{5}}{100} * \frac{100 - (100 - S_2) * \frac{F_2}{5}}{100} * \dots * \frac{100 - (100 - S_n) * \frac{F_n}{5}}{100}$$

[00292] where  $O$  is the overall compatibility score,  $n$  is the total number of rule sets 28 and workload types 30 included in the analysis,  $S_i$  is the compatibility score of the  $i^{\text{th}}$  rule set 28 or workload type 30 and  $F_i$  is the importance factor of the  $i^{\text{th}}$  rule set 28 or workload type 30.

[00293] It can be appreciated that setting the importance factor 88 to zero eliminates the contribution of the corresponding score to the overall score. Also, setting the importance factor to a value less than 5 reduces the score penalty by 20% to %100 of its original value.

[00294] For example, a compatibility score of 90 implies a score penalty of 10 (i.e.  $100 - 90 = 10$ ). Given an importance factor of 1, the adjusted score is 98 (i.e.  $100 - 10 * 1/5 = 100 - 2 = 98$ ). On the other hand, setting the importance factor to a value greater than 5 increases the score penalty by 20% to 100% of its original value. Using the above example, given a score of 90 and an importance factor of 10, the adjusted score would be 80 (i.e.  $100 - 10 * 10/5 = 100 - 20 = 80$ ). The range of importance factors 88 and their impact on the penalty scores are summarized below in Table 13.

Importance Factor	Affect on Score Penalty	Original Score	Original Score Penalty	Adjusted Penalty Score	Adjusted Score
0	-100%	90	10	0	100
1	-80%	90	10	2	98
2	-60%	90	10	4	96
3	-40%	90	10	6	94
4	-20%	90	10	8	92
5	0	90	10	10	90
6	+20%	90	10	12	88
7	+40%	90	10	14	86
8	+60%	90	10	16	84
9	+80%	90	10	18	82
10	+100%	90	10	20	80

*Table 13: Importance Factors*

[00295] If more systems 16 are to be examined, the above process is repeated. When overall compatibility analysis scores for all server pairs have been computed, the map 36 is displayed graphically (see Figure 6) and each cell 92 is linked to a scorecard 98 that provides further information (Figure 23). The further information can be viewed by selecting the cell 92. A sorting algorithm is then preferably executed to configure the map 36 as shown in Figure 6.

## **Visualization and Mapping of Compatibility Scores**

[00296] As mentioned above, the 1-to-1 compatibility analyses of N system computes NxN compatibility scores by individually considering each system 16 as a consolidation source and as a target. Preferably, the scores range from 0 to 100 with higher scores indicating greater system compatibility. The analysis will thus also consider the trivial cases where systems 16 are consolidated with themselves and would be given a maximum score, e.g. 100. For display and reporting purposes, the scores are preferably arranged in an NxN map form.

## **Rule-based Compatibility Analysis Visualization**

[00297] An example of a rule-based compatibility analysis map 32 is shown in Figure 4. The compatibility analysis map 32 provides an organized graphical mapping of system compatibility for each source/target system pair on the basis of configuration data. The map 32 shown in

Figure 4 is structured having each system 16 in the environment 12 listed both down the leftmost column and along the uppermost row. Each row represents a consolidation source system, and each column represents the possible consolidation target. Each cell 92 contains the score corresponding to the case where the row system is consolidated onto the column (target) system 16.

**[00298]** The preferred output shown in Figure 4 arranges the systems 16 in the map 32 such that a 100% compatibility exists along the diagonal where each server is naturally 100% compatible with itself. The map 32 is preferably displayed such that each cell 92 includes a numerical score and a shade of a certain colour. As noted above, the higher the score (from zero (0) to one hundred (100)), the higher the compatibility. The scores are pre-classified into predefined ranges that indicate the level of compatibility between two systems 16. Each range maps to a corresponding colour or shade for display in the map 32. For example, the following ranges and colour codes can be used: score = 100, 100% compatible, dark green; score = 75-99, highly compatible, green; score = 50-74, somewhat compatible, yellow; score = 25-49, low compatibility, orange; and score = 0-24, incompatible, red.

**[00299]** The above ranges are only one example. Preferably, the ranges can be adjusted to reflect more conservative and less conservative views on the compatibility results. The ranges can be adjusted using a graphical tool similar to a contrast slider used in graphics programs. Adjustment of the slider would correspondingly adjust the ranges and in turn the colours. This allows the results to be tailored to a specific situation.

**[00300]** It is therefore seen that the graphical output of the map 32 provides an intuitive mapping between the source/target pairs in the environment 12 to assist in visualizing where compatibilities exist and do not exist. In Figure 4, it can be seen that a system pair having a score = 100 indicates complete compatibility between the two systems 16 for the particular strategy being observed, e.g. based on a chosen rule set(s) 28. It can also be seen that a system pair with a relatively lower score such as 26 is relatively less compatible for the strategy being observed.

1 [00301] The detailed differences shown in Figures 19 can be viewed by clicking on the  
2 relevant cell 92. Selecting a particular cell 92 accesses the detailed differences table 100 shown  
3 in Figure 19 which shows the important differences between the two systems, the rules and  
4 weights that were applied and preferably a remediation cost for making the servers more  
5 compatible. As shown in Figure 19, a summary differences table 102 may also be presented  
6 when selecting a particular cell 92, which lists the description of the differences and the weight  
7 applied for each difference, to give a high level overview of where the differences arise.

## 8 **System Workload Compatibility Visualization**

9 [00302] An example workload compatibility analysis map 34 is shown in Figure 5. The map  
10 34 is the analog of the map 32 for workload analyses. The map 34 includes a similar graphical  
11 display that indicates a score and a colour or shading for each cell to provide an intuitive  
12 mapping between candidate source/target server pairs. The workload data is obtained using tools  
13 such as the table 76 shown in Figure 8 and corresponds to a particular workload factor, e.g. CPU  
14 utilization, network I/O, disk I/O, etc. A high workload score indicates that the candidate server  
15 pair being considered has a high compatibility for accommodating the workload on the target  
16 system. The specific algorithms used in determining the score are discussed in greater detail  
17 below. The servers are listed in the upper row and leftmost column and each cell 92 represents  
18 the compatibility of its corresponding server pair in the map. It can be appreciated that a  
19 relatively high score in a particular cell 92 indicates a high workload compatibility for  
20 consolidating to the target server, and likewise, relatively lower scores, e.g. 42 indicate lower  
21 workload compatibility for a particular system pair.

22 [00303] The workload analysis details shown in Figures 22 can be viewed by clicking on the  
23 relevant cell 92. Selecting a particular cell 92 accesses the information about the workload  
24 analysis that generated the score shown in Figure 22, which shows the key stacked workload  
25 values, the workload benchmarks that were applied and preferably workload charts for each  
26 system separately, and stacked together.

## Overall Compatibility Visualization

[00304] An example overall compatibility analysis map 36 is shown in Figure 6. The map 36 comprises a similar arrangement as the maps 32 and 34, which lists the servers in the uppermost row and leftmost column to provide 100% compatibility along the diagonal. Preferably the same scoring and shading convention is used by all types of compatibility maps. The map 36 provides a visual display of scoring for candidate system pairs that considers the rule-based compatibility maps 32 and the workload compatibility maps 34.

[00305] The score provided in each cell 92 indicates the overall compatibility for consolidating systems 16. It should be noted that in some cases two systems 16 can have a high configuration compatibility but a low workload compatibility and thus end up with a reduced or relatively low overall score. It is therefore seen that the map 36 provides a comprehensive score that considers not only the compatibility of systems 28 at the setting level but also in its utilization. By displaying the configuration maps 32, business maps, workload maps 34 and overall map 36 in a consolidation roadmap, a complete picture of the entire system can be ascertained in an organized manner. The maps 32, 34 and 36 provide a visual representation of the compatibilities and provide an intuitive way to evaluate the likelihood that systems can be consolidated, to analyse compliance and drive remediation measures to modify systems 16 so that they can become more compatible with other systems 16 in the environment 12. It can therefore be seen that a significant amount of quantitative data can be analysed in a convenient manner using the graphical maps 32, 34 and 36, and associated reports and graphs (described below).

[00306] For example, a system pair that is not compatible only for the reason that certain critical software upgrades have not been implemented, the information can be uncovered by the map 32, and then investigated, so that upgrades can be implemented, referred to herein as remediation. Remediation can be determined by modeling cost of implementing upgrades, fixes etc that are needed in the rule sets. If remediation is then implemented, a subsequent analysis may then show the same server pair to be highly compatible and thus suitable candidates for consolidation.

1 [00307] The overall analysis details 98 shown in Figures 23 can be viewed by clicking on the  
2 relevant cell 92. Selecting a particular cell 92 accesses the information about the rule-based and  
3 workload analyses that generated the score shown in Figure 23, which shows the key differences  
4 and stacked workload values and charts.

## 5 **Sorting Examples**

6 [00308] The maps 32, 34 and 36 can be sorted in various ways to convey different  
7 information. For example, sorting algorithms such as a simple row sort, a simple column sort  
8 and a sorting by group can be used.

9 [00309] A simple row sort involves computing the total scores for each source system (by  
10 row), and subsequently sorting the rows by ascending total scores. In this arrangement, the  
11 highest total scores are indicative of source systems that are the best candidates to consolidate  
12 onto other systems.

13 [00310] A simple column sort involves computing the total scores for each target system (by  
14 column) and subsequently sorting the columns by ascending total score. In this arrangement, the  
15 highest total scores are indicative of the best consolidation target systems.

16 [00311] Sorting by group involves computing the difference between each system pair, and  
17 arranging the systems to minimize the total difference between each pair of adjacent systems in  
18 the map. The difference between a system pair can be computed by taking the square root of the  
19 sum of the squares of the difference of a pair's individual compatibility score against each other  
20 system in the analysis. In general, the smaller the total difference between two systems, the  
21 more similar the two systems with respect to their compatibility with the other systems. The  
22 group sort promotes the visualization of the logical breakdown of an environment by producing  
23 clusters of compatible systems 18 around the map diagonal. These clusters are indicative of  
24 compatible regions in the environment 12. In virtualization analysis, these are often referred to  
25 as "affinity regions."

## 26 **Analysis Results-User Interaction**



1 [00312] It can also be seen that users can customize and interact with the analysis program 10  
2 during the analysis procedure to sort map scores, modify colour coding (as discussed above),  
3 show/specify source/targets, adjust weights and limits etc., and to show workload charts. This  
4 interaction enables the user to modify certain parameters in the analysis to take into account  
5 differences in objectives and different environments to provide a more accurate analysis.

## 6 **Multi-Dimensional Compatibility Analysis**

7 [00313] The high level process flow of the multi-dimensional compatibility analysis is  
8 illustrated in Figure 24(a). In addition to the common compatibility analysis input, this analysis  
9 takes a consolidation solution as input. In contrast to the 1-to-1 compatibility analysis that  
10 evaluates the compatibility of each system pair, this multi-dimensional compatibility analysis  
11 evaluates the compatibility of each transfer set 23 specified in the consolidation solution.

12 [00314] The multi-dimensional compatibility analysis extends the original 1-to-1  
13 compatibility analysis that assessed the transfer of a single source entity to a target. As with the  
14 1-to-1 compatibility analysis, the multi-dimensional analysis produces an overall compatibility  
15 scorecard 98 based on technical, business and workload constraints. Technical and business  
16 compatibility are evaluated through one or more rule sets 28. Workload compatibility is  
17 assessed through one or more workload types 30.

18 [00315] This produces multi-dimensional compatibility analysis results, which includes multi-  
19 dimensional compatibility scores, maps and details based on the proposed transfer sets 23.

20 [00316] For each transfer set 23, a compatibility score is computed for each rule set 28 and  
21 workload type 30. An overall compatibility score the transfer set 23 is then derived from the  
22 individual scores. For example, consider an analysis comprised of 20 systems, 3 rule sets, 2  
23 workload types and 5 transfer sets 23:

- 24 • Systems: S1, S2, S3, ... S20
- 25 • Analyzed with rule sets: R1, R2, R3
- 26 • Analyzed with workload types: W1, W2

- 1       • Transfer sets:
  - 2           ○ T1 (S1, S2, S3 stacked onto S4)
  - 3           ○ T2 (S5, S6, S7, S8, S9 stacked onto S10)
  - 4           ○ T3 (S11, S12, S13 stacked onto S14)
  - 5           ○ T4 (S15 stacked onto S16)
  - 6           ○ T5 (S17 stacked onto S18)
- 7       • Unaffected systems: S19, S20

8   **[00317]** For the above example, the multi-dimensional compatibility analysis would comprise  
9   5 overall compatibility scores (one for each transfer set), 15 rule-based compatibility scores (5  
10 transfer sets x 3 rule sets), and 10 workload compatibility scores (5 transfer sets x 2 workload  
11 types).

12 **[00318]** The systems 16 referenced in the transfer sets 23 of the consolidation solution  
13 correspond to the systems 16 specified in the analysis input. Typically, the consolidation  
14 solution is manually specified by the user, but may also be based on the consolidation analysis,  
15 as described later.

16 **[00319]** In addition to evaluating the compatibility of the specified transfer sets, the  
17 compatibility analysis can evaluate the incremental effect of adding other source systems  
18 (specified in the analysis input) to the specified transfer sets. From the above example consisting  
19 of systems S1 to S20, the compatibility of the source systems S5 to S20 can be individually  
20 assessed against the transfer set T1. Similarly, the compatibility of the source systems S1 to S4  
21 and S11 to S20 can be assessed with respect to the transfer set T2.

22 **[00320]** Similar to the 1-to-1 compatibility analysis, this analysis involves 4 stages. The first  
23 stage is gets the system data 18 required for the analysis to produce the analysis data snapshot.  
24 The second stage performs a multi-dimensional compatibility analysis for each rule set 28 for  
25 each transfer set 23. Next, the workload compatibility analysis is performed for each workload  
26 type 30 for each transfer set 23. Finally, these analysis results are combined to determine overall  
27 compatibility of each transfer set.

1 [00321] The multi-dimensional rule-based compatibility analysis differs from the 1-to-1  
2 compatibility analysis since a transfer set can include multiple sources (N) to be transferred to  
3 the target, the analysis may evaluate the compatibility of sources amongst each other (N-by-N)  
4 as well as each source against the target (N-to-1) as will be explained in greater detail below.  
5 The multi-dimensional workload and overall compatibility analysis algorithms are analogous to  
6 their 1-to-1 analysis counterparts.

## 7 **Multi-dimensional Rule-based Compatibility Analysis**

8 [00322] To assess the compatibility of transferring multiple source entities (N) to a target (1),  
9 the rule-based analysis can compute a compatibility score based on a combination of N-to-1 and  
10 N-by-N compatibility analyses. An *N-to-1 intercompatibility* analysis assesses each source  
11 system against the target. An *N-by-N intracompatibility* analysis evaluates each source system  
12 against each of the other source systems. This is illustrated in a process flow diagram in Figure  
13 24(b).

14 [00323] Criteria used to choose when to employ an N-to-1, N-by-N or both compatibility  
15 analyses depend upon the target type (concrete or malleable), consolidation strategy (stacking or  
16 virtualization), and nature of the rule item.

17 [00324] Concrete target models are assumed to rigid with respect to their configurations and  
18 attributes such that source entities to be consolidated are assumed to be required to conform to  
19 the target. To assess transferring source entities onto a concrete target, the N-to-1 inter-  
20 compatibility analysis is performed. Alternatively, malleable target models are generally  
21 adaptable in accommodating source entities to be consolidated. To assess transferring source  
22 entities onto a malleable target, the N-to-1 inter-compatibility analysis can be limited to the  
23 aspects that are not malleable.

24 [00325] When *stacking* multiple source entities onto a target, the source entities and targets  
25 coexist in the same operating system environment. Because of this inherent sharing, there is  
26 little flexibility in accommodating individual application requirements, and thus the target is  
27 deemed to be concrete. As such, the multi-dimensional analysis considers the N-to-1 inter-

1 compatibility between the source entities and the target as the primary analysis mechanism, but,  
2 depending on the rule sets in use, may also consider the N-by-N intra-compatibility of the source  
3 entities amongst each other.

4 **[00326]** When virtualizing multiple source entities onto a target, the source entities are often  
5 transferred as separate virtual images that run on the target. This means that there is high  
6 isolation between operating system-level parameters, and causes virtualization rule sets to  
7 generally ignore such items. What is relevant, however, is the affinity between systems at the  
8 hardware, storage and network level, and it is critical to ensure that the systems being combined  
9 are consistent in this regard. In general, this causes the multi-dimensional analysis to focus on  
10 the N-to-N compatibility within the source entities, although certain concrete aspects of the  
11 target systems (such as processor architecture) may still be subjected to (N-to-1) analysis.

## 12 **N-to-1 Intercompatibility Score Calculation**

14 **[00327]** N-to-1 intercompatibility scores reflect the compatibility between N source entities  
15 and a single target as defined by a transfer set 23 as shown in Figure 24(c). This analysis is  
16 performed with respect to a given rule set and involves: 1) Separately evaluate each source entity  
17 against the target with the rule set to compile a list of the union of all matched rule items; 2) For  
18 each matched rule item, use the rule item's mutex (mutually exclusive) flag to determine whether  
19 to count duplicate matched rule items once or multiple times; and 3) Compute the score based on  
20 the product of all the penalty weights associated with the valid matched rule items:

21 **[00328]** 
$$S = 100 * (1 - w_1) * (1 - w_2) * (1 - w_3) * \dots (1 - w_n);$$

22 **[00329]** where S is the score and  $w_i$  is the penalty weight of the  $i^{\text{th}}$  matched item.

## 23 **N-to-1 Score Example**

25 **[00330]** For example, assuming a transfer set t1 comprises of systems s1, s2 and s3 stacked  
26 onto s16, the union of matched rule items is based on evaluating s1 against s16, s2 against s16  
27 and s3 against s16. Assuming this analysis produces a list of matched items comprising those  
28 shown below in Table 14:

21635070.1

#	Source	Target	Rule Item	Source Value	Target Value	Mutex	Weight
1	S1	S16	Different Patch Levels	SP2	SP3	Y	0.03
2	S1	S16	Different Default Gateways	10.0.0.1	192.168.0.1	N	0.02
3	S2	S16	Different Patch Levels	SP2	SP3	Y	0.03
4	S3	S16	Different Patch Levels	SP4	SP3	Y	0.03
5	S3	S16	Different Default Gateways	10.0.0.1	192.168.0.1	N	0.02
6	S3	S16	Different Boot Settings	TRUE	FALSE	N	0.01

*Table 14: Matched Items – Multi-Dimensional N-to-1 example*

**[00331]** Although the target and source values vary, items 1, 3 and 4 apply to the same rule item and are treated as duplicates due to the enabled mutex flag so that the penalty weight is applied only once. Items 2 and 5 apply to the same item and are exact duplicates (same values) so the penalty weight is applied only once, even though the mutex flag is not enabled for this item. As such, the compatibility score is computed as follows:

**[00332]**  $S = 100 * (1 - 0.03) * (1 - 0.02) * (1 - 0.01) = 94$

#### **N-by-N Intracompatibility Score Calculation**

**[00333]** N-by-N intracompatibility scores reflect the compatibility amongst N source entities with respect to a given rule set as shown in Figure 24(d). This analysis involves: 1) Separately evaluate each source entity against the other source entities with the rule set to compile a list of the union of all matched rule items; 2) For each matched rule item, use the rule item's mutex (mutually exclusive) flag to determine whether to count duplicate matched rule items once or multiple times; and 3) Compute the score based on the product of all the penalty weights associated with the valid matched rule items:

**[00334]**  $S = 100 * (1 - w_1) * (1 - w_2) * (1 - w_3) * \dots * (1 - w_n);$

**[00335]** where S is the score and  $w_i$  is the penalty weight of the  $i^{th}$  matched item.

#### **N-by-N Score Example**

21635070.1

[00336] For example, assuming a transfer set t1 comprises of systems s1, s2 and s3 stacked onto s16, the union of matched rule items is based on evaluating s1 against s2, s2 against s1, s2 against s3, s3 against s2, s1 against s3 and s3 against s1. Assuming this analysis produces a list of matched items comprising those shown below in Table 15:

#	Source	Target	Rule Item	Source Value	Target Value	Mutex	Weight
1	S1	S3	Different Patch Levels	SP2	SP4	Y	<b>0.03</b>
2	S3	S1	Different Patch Levels	SP4	SP2	Y	0.03
3	S2	S3	Different Patch Level	SP2	SP4	Y	0.03
4	S3	S2	Different Patch Level	SP4	SP2	Y	0.03
5	S1	S2	Different Default Gateways	10.0.0.1	192.168.0.1	N	<b>0.02</b>
6	S2	S1	Different Default Gateways	192.168.0.1	10.0.0.1	N	0.02
7	S3	S2	Different Default Gateways	10.0.0.1	192.168.0.1	N	0.02
8	S2	S3	Different Default Gateways	192.168.0.1	10.0.0.1	N	0.02
9	S1	S3	Different Boot Settings	TRUE	FALSE	N	<b>0.01</b>
10	S3	S1	Different Boot Settings	FALSE	TRUE	N	0.01
11	S2	S3	Different Boot Settings	TRUE	FALSE	N	0.01
12	S3	S2	Different Boot Settings	FALSE	TRUE	N	0.01

*Table 15: Matched Items – Multi-Dimensional N-by-N example*

[00337] Items 1-4, 5-8 and 9-12, respectively are duplicates as they apply to the same rule items and have the same values. The compatibility score is computed as follows:

[00338]  $S = 100 * (1 - 0.03) * (1 - 0.02) * (1 - 0.01) = 94.$

### Multi-dimensional Workload Compatibility Analysis

[00339] A procedure for stacking the workload of multiple source systems on a target system is shown in Figure 25. The multi-stacking procedure considers the workload limits that is

21635070.1

specified using the program 150, the per-system workload benchmarks (e.g. CPU power), and the data snapshot containing the workload data for the source and target systems 16 that comprise the transfer sets 23 to analyze. The analysis may evaluate transfer sets 23 with any number of sources stacked on a target for more than one workload type 30.

[00340] For each workload type 30, each transfer set 23 is evaluated. For each source in the transfer set 23, the system benchmarks are used to normalize the workloads as discussed above, and the source workload is stacked on the target system. Once every source in the set is stacked on the target system, the workload compatibility score is computed as discussed above. The above is repeated for each transfer set 23. A multi-stack report may then be generated, which gives a workload compatibility scorecard for the transfer sets along with workload compatibility scoring details and normalized multi-stacked workload charts.

[00341] Sample multi-dimensional compatibility analysis results are shown in Figure 26-28. Figure 26 shows a compatibility score map 110 with the analyzed transfer sets. Figures 27 and 28 show the summary 112 and charts 114 from the multi-stack workload compatibility analysis details.

## **Consolidation Analysis**

[00342] The consolidation analysis process flow is illustrated as D in Figure 9. Using the common compatibility analysis input and additional auto fit inputs, this analysis seeks the consolidation solution that maximizes the number of transfers while still fulfilling the several pre-defined constraints. The consolidation analysis repeatedly employs the multi-dimensional compatibility analysis to assess potential transfer set candidates. The result of the consolidation analysis comprises of the consolidation solution and the corresponding multi-dimensional compatibility analysis.

[00343] A process flow of the consolidation analysis is shown in Figure 29.

[00344] The auto fit input includes the following parameters: transfer type (e.g. virtualize or stacking), minimum allowable overall compatibility score for proposed transfer sets, minimum

1 number of source entities to transfer per target, maximum number of source entities to transfer  
2 per target, and quick vs. detailed search for the best fit. Target systems can also be designated as  
3 malleable or concrete models.

4 **[00345]** As part of a compatibility analysis input specification, systems can be designated for  
5 consideration as a source only, as a target only or as either a source or a target. These  
6 designations serve as constraints when defining transfers in the context of a compatibility  
7 analysis. The analysis can be performed on an analysis with pre-existing source-target transfers.  
8 Analyses containing systems designated as source or target-only (and no source or target  
9 designations) are referred to as "directed analysis."

10 **[00346]** The same transfer type may be assumed for all automatically determined transfers  
11 within an analysis. The selected transfer type affects how the compatibility analysis is  
12 performed. The minimum overall compatibility score dictates the lowest allowable score  
13 (sensitivity) for the transfer sets to be included in the consolidation solution. Lowering the  
14 minimum allowable score permits a greater degree of consolidation and potentially more  
15 transfers. The minimum and maximum limits for source entities to be transferred per target  
16 (cardinality) define additional constraints on the consolidation solution. The quick search  
17 performs a simplified form of the auto fit calculation, whereas the detailed search performs a  
18 more exhaustive search for the optimal solution. This distinction is provided for quick  
19 assessments of analyses containing a large numbers of systems to be analyzed.

20 **[00347]** The transfer auto fit problem can be considered as a significantly more complex form  
21 of the classic bin packing problem. The bin packing problem involves packing objects of  
22 different volumes into a finite number of bins of varying volumes in a way that minimizes the  
23 number of bins used. The transfer auto fit problem involves transferring source entities onto a  
24 finite number of targets in a way that maximizes the number of transfers. The basis by which  
25 source entities are assessed to "fit" onto targets is based on the highly nonlinear compatibility  
26 scores of the transfer sets. As a further consideration, which can increase complexity, some  
27 entities may be either source or targets. The auto fit problem is a combinatorial optimization  
28 problem that is computationally expensive to solve through a brute force search of all possible



transfer set permutations. Although straightforward to implement, this exhaustive algorithm is impractical due to its excessive computational and resource requirements for medium to large data sets. Consequently, this class of problem is most efficiently solved through heuristic algorithms that yield good but likely suboptimal solutions.

[00348] There are four variants of the heuristic auto fit algorithm that searches for the best consolidation solution:

[00349] Quick Stack – quick search for a stacking-based consolidation solution;

[00350] Detailed Stack – more comprehensive search for a stacking-based consolidation solution;

[00351] Quick Virtualization – quick search for a virtualization-based consolidation solution; and

[00352] Detailed Virtualization – more comprehensive search for a virtualization-based consolidation solution.

[00353] The auto fit algorithms are iterative and involve the following common phases:

#### **Compile Valid Source and Target Candidates**

[00354] The initial phase filters the source and target lists by eliminating invalid entity combinations based on the 1-to-1 compatibility scores that are less than the minimum allowable compatibility score. It also filters out entity combinations based on the source-only or target-only designations.

#### **Set up Auto Fit Parameters**

[00355] The auto fit algorithm search parameters are then set up. The parameters can vary for each algorithm. Example search parameters include the order by which sources and targets are processed and the criteria for choosing the best transfer set.

#### **Compile Candidate Transfer Sets**

21635070.1

1 [00356] The next phase compiles a collection of candidate transfer sets 23 from the available  
2 pool of sources and targets. The candidate transfer sets 23 fulfill the auto fit constraints (e.g.  
3 minimum allowable score, minimum transfers per transfer set, maximum transfers per transfer  
4 set). The collection of candidate transfer sets may not represent a consolidation solution (i.e.  
5 referenced sources and targets may not be mutually exclusive amongst transfer sets 23). The  
6 algorithms vary in the criteria employed in composing the transfer sets. In general, the detailed  
7 search algorithms generate more candidate transfer sets than quick searches in order to assess  
8 more transfer permutations.

#### 9 **Choose Best Candidate Transfer Set**

10  
11 [00357] The next phase compares the candidate transfer sets 23 and chooses the “best”  
12 transfer set 23 amongst the candidates. The criteria employed to select the best transfer set 23  
13 varies amongst the algorithms. Possible criteria include the number of transfers, the  
14 compatibility score, general compatibility of entities referenced by set and whether the transfer  
15 set target is a target-only.

#### 16 **Add Transfer Set to Consolidation Solution**

17 [00358] Once a transfer set is chosen, it is added to the intermediate consolidation solution.  
18 The entities referenced by the transfer set are removed from the list of available sources and  
19 targets and the three preceding phases are repeated until the available sources or targets are  
20 consumed.

#### 21 **Compile Consolidation Solution Candidates**

22  
23 [00359] Once all the sources or targets are consumed or ruled out, the consolidation solution  
24 is considered complete and added to a list of candidate solutions. Additional consolidation  
25 solutions can be compiled by iterating from the second phase with variations to the auto fit  
26 parameters for compiling and choosing candidate transfer sets.

#### 27 **Choose Best Consolidation Solution**

[00360] The criteria used to stop compiling additional solutions can be based on detecting that the solution is converging on a pre-defined maximum number of iterations.

[00361] Finally, the best candidate consolidation solution can be selected based on some criteria such as the largest reduction of systems with the highest average transfer set scores. The general algorithm is shown in the flow diagram depicted in Figure 30.

## Auto Fit Example

[00362] The following example demonstrates how a variant of the auto fit algorithm searches for a consolidation solution for a compatibility analysis comprised of 20 systems (S1, S2, S3, ... S20) where 15 systems (S1-15) have been designated to be source-only and 5 (S16-20) to be target-only. For this example, the auto fit input parameters are: 1) Transfer type: stacking; 2) Minimum allowable compatibility score: 80; 3) Minimum sources per target: 1; 4) Maximum sources per target: 5; and 6) Search type: quick.

[00363] The auto fit would be performed as follows:

### 1 – Compile Valid Source and Target Candidates

[00364] For each target (S16-S20), compile the list of possible sources. Since some of the 1-to-1 source-target compatibility scores are less than 80 (specified minimum allowed score), the following source-target candidates are found as shown in Table 16.

Target	Sources (1-to-1 scores)
S16	S1 (100), S2 (95), S3 (90), S4 (88), S5 (88), S6 (85), S7 (82), S8 (81)
S17	S1, S2, S3, S5, S6, S7, S8, S9, S10
S18	S1, S2, S3, S6, S7, S8, S9, S10, S11
S19	S9, S10, S11, S12, S13, S14, S15
S20	S9, S10, S11, S12, S13, S14, S15

Table 16: Target-Source Candidates

### 2 – Setup Auto Fit Parameters

[00365] The auto fit search parameters initialized for this iteration assume the following: 1) When compiling candidate transfer sets 23, sort source systems in descending order when

1 stacking onto target; and 2) When choosing the best transfer set 23, choose set with most  
2 transfers and if there is a tie, choose the set 23 with the higher score.

### 3 **3 – Compile Candidate Transfer Sets**

4  
5 **[00366]** The candidate transfer sets 23 are then compiled from the source-target candidates.  
6 For each target, the candidate sources are sorted in descending order and are incrementally  
7 stacked onto the target. The transfer set score is computed as each source is stacked and if the  
8 score is above the minimum allowable score, the source is added to the transfer set 23. If the  
9 score is below the minimum allowable, the source is not included in the set. This process is  
10 repeated until all the sources have been attempted or the maximum number of sources per target  
11 has been reached.

12 **[00367]** For this quick search algorithm, only a single pass is performed for each target so that  
13 only one transfer set candidate is created per target. Other search algorithms can perform  
14 multiple passes per target to assess more transfer permutations. The candidate transfer sets 23  
15 are shown below in Table 17.

Transfer Set	Target	Sources	# Sources	Score
T1	S16	S1, S2, S3, S4, S5	5	82
T2	S17	S1, S2, S3, S5	4	81
T3	S18	S1, S2, S3, S6, S7, S10	6	85
T4	S19	S9, S10, S11, S12, S13, S14	6	83
T5	S20	S9, S10, S11, S12, S13, S14	6	84

16 *Table 17: Candidate Transfer Sets*

### 18 **4 – Choose Best Candidate Transfer Set**

19  
20 **[00368]** The transfer set T3 is chosen as the best transfer set from the candidates. For this  
21 example, the criteria are the greatest number of sources and if tied, the highest score.

### 22 **5 – Add Transfer Set to Consolidation Solution**

23 **[00369]** The transfer set T3 is added to the consolidation solution, and the entities referenced  
24 by this transfer set are removed from the target-source candidates list. The updated target-source  
25 candidates are shown below in Table 18.

Target	Sources (1-to-1 scores)
S16	S4, S5, S8
S17	S5, S8, S9
S19	S9, S11, S12, S13, S14, S15
S20	S9, S11, S12, S13, S14, S15

*Table 18: Updated Target-Source Candidates*

[00370] Since there are available target-source candidates, another iteration to compile candidate transfer sets and choose the best set is performed. These iterations are repeated until there are no more target-source candidates, at which time a consolidation solution is considered complete. The consolidation solution candidates are shown below in Table 19.

Target	Sources	# Sources	Score
S16	S4, S5, S8	3	86
S18	S1, S2, S3, S6, S7, S10	6	85
S19	S9, S11, S12, S15	6	83
S20	S13, S14	2	87

*Table 19: Consolidation Solution Candidates*

## 6 – Compile Consolidation Solution Candidates

[00371] The consolidation solution is then saved, and if warranted by the selected algorithm, another auto fit can be performed with a different search parameters (e.g. sort source systems in ascending order before stacking onto target) to generate another consolidation solution.

## 7 – Choose Best Consolidation Solution

[00372] The consolidation solution candidates are compared and the best one is chosen based on some pre-defined criteria. A sample consolidation solution summary 116 is shown in Figure 31.

## Example Compatibility Analysis

[00373] Compatibility and consolidation analyses are described by way of example only below to illustrate an end-to-end data flow in conducting complete analyses for an arbitrary environment 12. These analyses are performed through the web client user interface 74. These

examples assume that the requisite system data 18 for the analyses have already been collected and loaded into the data repository 54 and caches 56 and 58.

### **1-to-1 Compatibility Analysis Example**

**[00374]** This type of analysis typically involves of the following steps: 1) Create a new analysis in the desired analysis folder; 2) Specify the mandatory analysis input; 3) Optionally, adjust analysis input parameters whose default values are not desired; 4) Run the analysis; and 5) View the analysis results.

**[00375]** The analysis can be created in an analysis folder on a computer 14, to help organize the analyses. Figure 32 shows an example analysis folder hierarchy containing existing analyses.

**[00376]** An analysis is created through a web page dialog as shown in the screen shot in Figure 33. The analysis name is preferably provided along with an optional description. Other analysis inputs comprise a list of systems 16 to analyze and one or more rule sets 28 and/or workload types 30 to apply to evaluate the 1-to-1 compatibility of the systems 16.

**[00377]** In this example, two rule sets 28 and one workload type 30 are selected. The following additional input may also be specified if the default values are not appropriate: 1) Adjustment of one or more rule weights, disable rules, or modify remediation costs in one or more of the selected rule sets; 2) Adjustment of the workload data date range; 3) Adjustment of one or more workload limits of the selected workload types; 4) Selection of different workload stacking and scoring parameters; and 5) Changing the importance factors for computing the overall scores. Figures 34, 35 and 36 show the pages used to customize rule sets 28, workloads 30 and importance factors 88, respectively. Once the analysis input is specified, the analysis can be executed. The analysis results can be viewed through the web client user interface 44/74. The results include the 1-to-1 compatibility maps for the overall and one for each rule set 28 and workload type 30. Such compatibility maps and corresponding analysis map details are shown in Figures 37 to 45.

## **Multi-dimensional Compatibility Analysis Example**

[00378] Continuing with the above example, one or more transfer sets 23 can be defined and the transfer sets 23 evaluated through a multi-dimensional compatibility analysis. The transfer sets 23 can be defined through the user interface, and are signified by one or more arrows that connect the source to the target. The color of the arrow can be used to indicate the transfer type (see Figure 46). Once the transfer sets 23 have been defined, the net effect mode may be selected to run the multi-dimensional analysis. In the resulting compatibility maps, the score in the cells that comprise the transfer set 23 reflects the multi-dimensional compatibility score (see Figure 47).

[00379] The corresponding overall compatibility details report is shown in Figure 48. Note that there are two sources transferred to the single target.

## **Consolidation Analysis Example**

[00380] Again continuing from the above example; a consolidation analysis may be performed to search for a consolidation solution by automating the analysis of the multi-dimensional scenarios. The input screen for the consolidation analysis is shown in Figure 49. Users can specify several parameters including the minimum allowable overall compatibility score and the transfer type. As well, users can choose to keep or remove existing transfer sets before performing the consolidation analysis.

[00381] Once specified, the consolidation analysis can be executed and the chosen transfer sets 23 are presented in the map as shown in Figure 50. The multi-dimensional compatibility score calculations are also used. Finally, a consolidation summary is provided as shown in Figures 51 and 52. Figure 51 shows that should the proposed transfers be applied, 28 systems would be consolidated to 17 systems, resulting in a 39% reduction of systems. Figure 52 lists the actual transfers proposed by the consolidation analysis.

## **Commentary**

[00382] Accordingly, the compatibility and consolidation analyses can be performed on a collection of system to 1) evaluate the 1-to-1 compatibility of every source-target pair, 2)

21635070.1

1 evaluate the multi-dimensional compatibility of specific transfer sets, and 3) to determine the  
2 best consolidation solution based on various constraints including the compatibility scores of the  
3 transfer sets. Though these analyses share many common elements, they can be performed  
4 independently.

5 **[00383]** These analyses are based on collected system data related to their technical  
6 configuration, business factors and workloads. Differential rule sets and workload compatibility  
7 algorithms are used to evaluate the compatibility of systems. The technical configuration,  
8 business and workload related compatibility results are combined to create an overall  
9 compatibility assessment. These results are visually represented using color coded scorecard  
10 maps.

11 **[00384]** It will be appreciated that although the system and workload analyses are performed  
12 in this example to contribute to the overall compatibility analyses, each analysis is suitable to be  
13 performed on its own and can be conducted separately for finer analyses. The finer analysis may  
14 be performed to focus on the remediation of only configuration settings at one time and  
15 spreading workload at another time. As such, each analysis and associated map may be  
16 generated on an individual basis without the need to perform the other analyses.

17 **[00385]** It will be appreciated that each analysis and associated map discussed above may  
18 instead be used for purposes other than consolidation such as capacity planning, regulatory  
19 compliance, change, inventory, optimization, administration etc. and any other purpose where  
20 compatibility of systems is useful for analyzing systems 16. It will also be appreciated that the  
21 program 10 may also be configured to allow user-entered attributes (e.g. location) that are not  
22 available via the auditing process and can factor such attributes into the rules and subsequent  
23 analysis.

24 **[00386]** It will further be appreciated that although the examples provided above are in the  
25 context of a distributed system of computer servers, the principles and algorithms discusses are  
26 applicable to any system having a plurality of sub-systems where the sub-systems perform  
27 similar tasks and thus are capable theoretically of being consolidation. For example, a local



1 network having a number of personal computers (PCs) could also benefit from a consolidation  
2 analysis.

3 [00387] Although the invention has been described with reference to certain specific  
4 embodiments, various modifications thereof will be apparent to those skilled in the art without  
5 departing from the spirit and scope of the invention as outlined in the claims appended hereto.

**Claims:**

1. A method for determining a consolidation solution for a plurality of computer systems comprising:
  - a) obtaining a data set comprising a compatibility score for each pair of said plurality of computer systems, each said compatibility score being indicative of the compatibility of one of said plurality of computer systems with respect to another of said plurality of computer systems;
  - b) determining one or more candidate transfer sets each indicating one or more of said computer systems capable of being transferred to a target computer system;
  - c) selecting a desired one of said one or more candidate transfer sets; and
  - d) providing said desired one as a consolidation solution.
2. The method according to claim 1 further comprising removing computer systems associated with said desired one of said one or more candidate transfer sets from said data set, repeating b) and c) to obtain additional consolidation solutions, and selecting a desired consolidation solution.
3. The method according to claim 1 comprising eliminating compatibility scores from said data set that do not meet a predefined threshold.
4. The method according to claim 3 wherein said predefined threshold is a minimum compatibility score.
5. The method according to claim 1 comprising defining one or more of a minimum transfers per transfer set and a maximum transfers per transfer set.
6. The method according to claim 1 wherein b) is performed according to a combinatorial optimization problem.

7. The method according to claim 6 wherein said optimization problem is solved through one or more heuristic algorithms.
8. The method according to claim 1 further comprising obtaining an indication of a transfer type, said transfer type being either virtualization or stacking.
9. The method according to claim 1 wherein said compatibility score is determined using a configuration compatibility score and a workload compatibility score.
10. A method for determining compatibilities for a plurality of computer systems comprising:
  - a) obtaining at least one transfer set indicating one or more of said computer systems capable of being transferred to a target computer system;
  - b) evaluating compatibilities of said one or more computer systems against said target computer system to obtain a first compatibility score;
  - c) evaluating compatibilities of each said one or more of said computer systems against each other to obtain a second compatibility score; and
  - d) computing an overall compatibility score for said transfer set using said first and second compatibility scores.
11. The method according to claim 10 wherein said first compatibility score is computed by separately evaluating each said one or more computer systems against said target computer system using a rule set defining one or more parameters indicative of compatibility.
12. The method according to claim 11 further comprising applying mutex flags to account for multiple instances of the same match.
13. The method according to claim 12 wherein said first compatibility score is computed by applying penalty weights associated with matched rule items and taking said mutex flags into consideration if necessary.

14. The method according to claim 10 wherein said second compatibility score is computed by separately evaluating each said one or more computer systems against said target computer system using a rule set defining one or more parameters indicative of compatibility.
15. The method according to claim 14 further comprising applying mutex flags to account for multiple instances of the same match.
16. The method according to claim 15 wherein said second compatibility score is computed by applying penalty weights associated with matched rule items and taking said mutex flags into consideration if necessary.
17. A computer program for determining compatibilities for a plurality of computer systems comprising an audit engine for obtaining information pertaining to the compatibility of said plurality of computer systems; an analysis engine for generating a compatibility score for each pair of said plurality of systems based on said information that is specific to respective pairs; and a client for displaying said compatibility score on an interface, said interface being configured to enable a user to specify parameters associated with a map summarizing said compatibility scores and to define and initiate a transfer of one or more of said computer systems to a target computer system.
18. A computer program product comprising computer executable instructions for causing a computer to perform the steps of claim 1.
19. A computer program product comprising computer executable instructions for causing a computer to perform the steps of claim 10.

1 ABSTRACT

2  
3 Compatibility and consolidation analyses can be performed on a collection of systems to  
4 evaluate the 1-to-1 compatibility of every source-target pair, evaluate the multi-dimensional  
5 compatibility of specific transfer sets, and to determine the best consolidation solution based on  
6 various constraints including the compatibility scores of the transfer sets. The analyses can be  
7 done together or be performed independently. These analyses are based on collected system data  
8 related to their technical configuration, business factors and workloads. Differential rule sets and  
9 workload compatibility algorithms are used to evaluate the compatibility of systems. The  
10 technical configuration, business and workload related compatibility results are combined to  
11 create an overall compatibility assessment. These results are visually represented using color  
12 coded scorecard maps




Figure 1

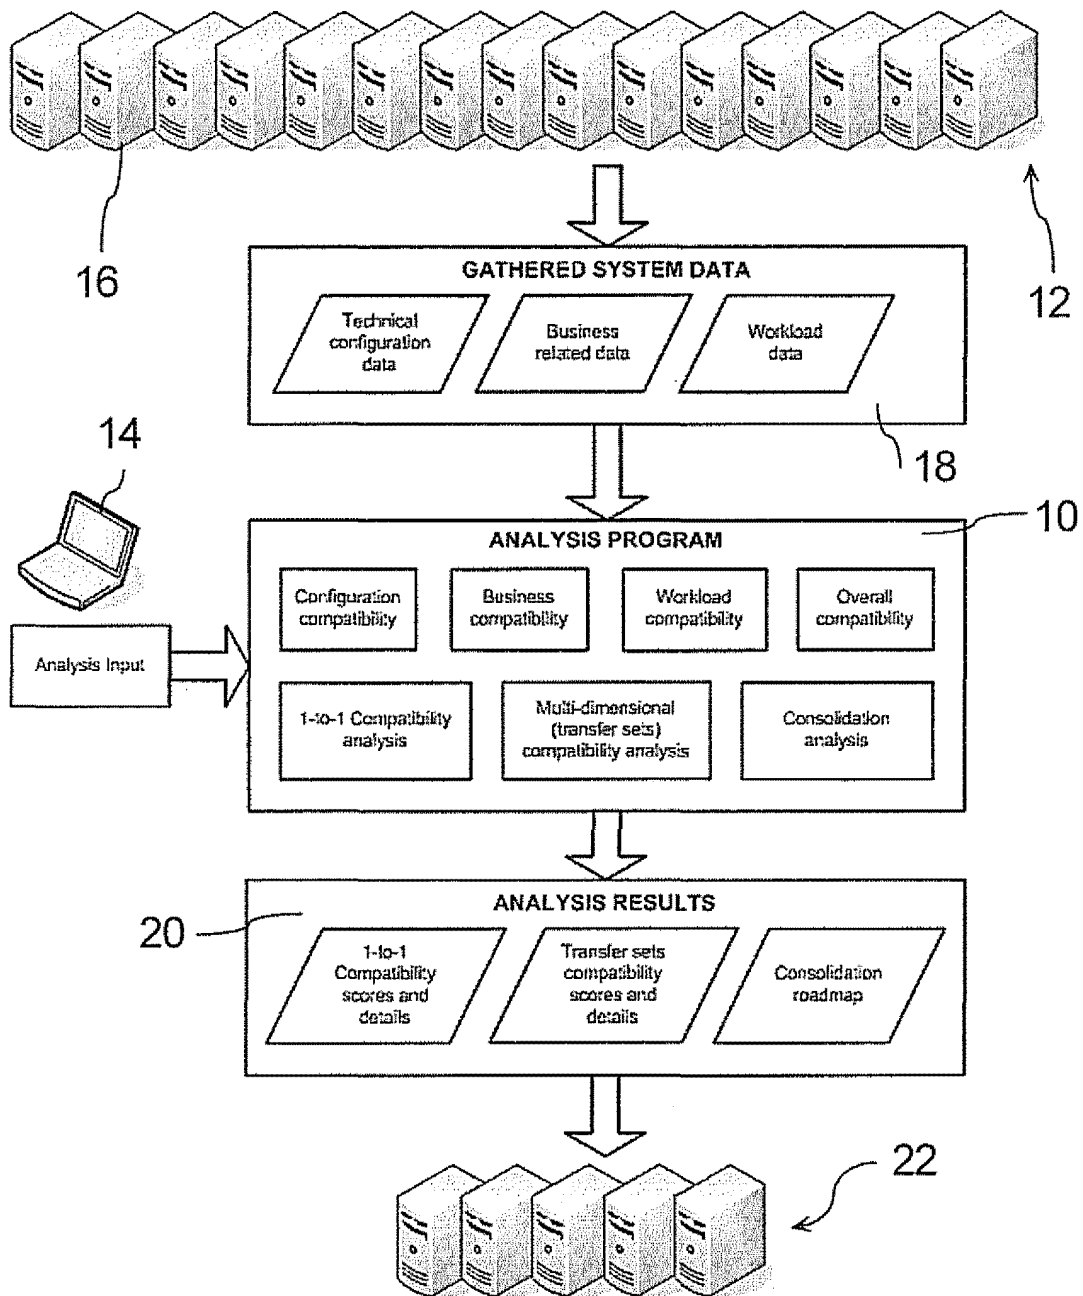


Figure 2




Figure 3




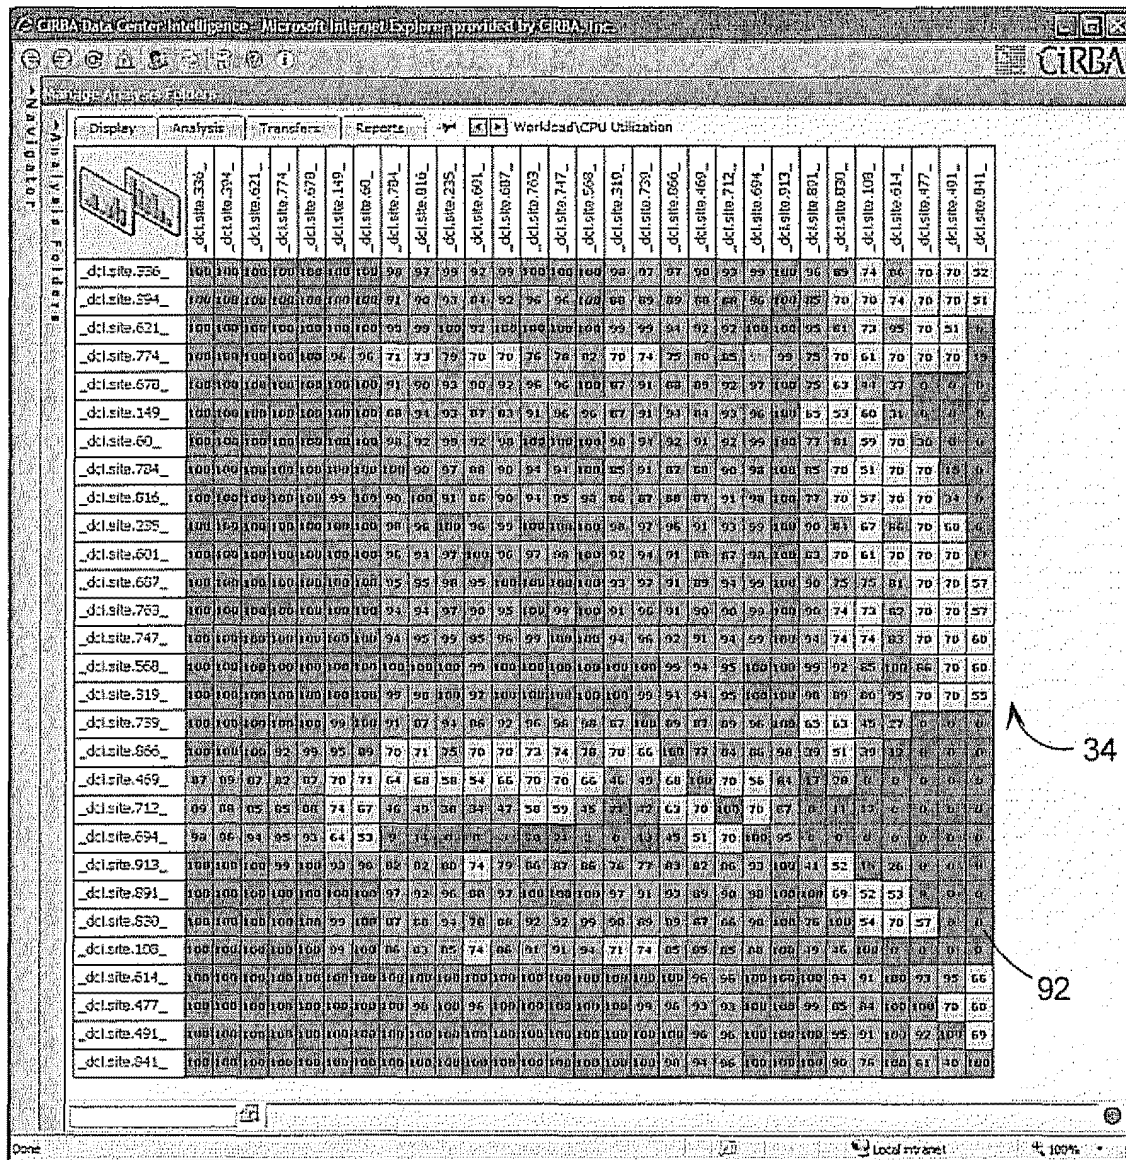


Figure 4



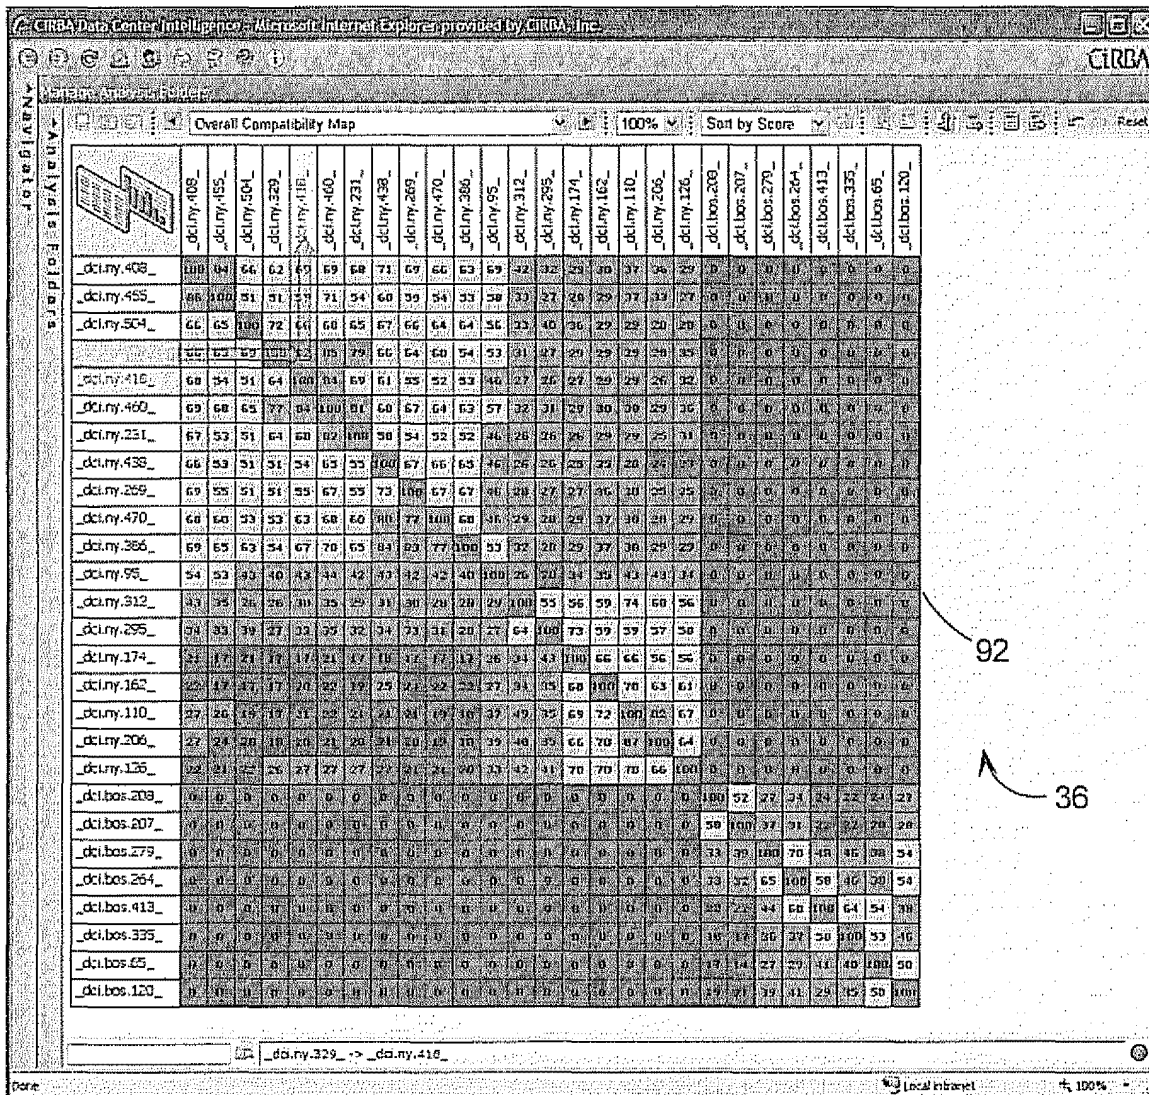


Figure 6




Figure 7

	Configuration Sources (SCI Analysis)										Workload Sources (WCI)				
	CIRBA Agentless					Third-Party Agents (Direct Audit)					CIRBA Agent		Data Import		
	SNMP	SSH	WMI	SQL	JMX	Compaq ILM	Dell OM	IBM Director	Concord, etc.	VMware ESX	Tivoli TME	BMC	SunMC	TADDM	Tivoli CM
Consolidation & Virtualization Analysis Strategy Instrumentation Enablement	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Database Consolidation	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Oracle database stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Oracle table stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Sybase database stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Sybase table stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
SQL Server - database stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
SQL Server - table stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
OS-Level Stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
UNIX Binary Compatibility	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Linux Binary Compatibility	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
UNIX Apps with file-based configurations	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
UNIX Apps with database-resident configurations	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Windows Binary Compatibility	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Windows Apps with file-based configurations	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Windows Apps with WMI-based configurations	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Windows Apps with database-resident configurations	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Application Server Stacking	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
WebLogic	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Webphere	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
IIS	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Jboss	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Apache Tomcat	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Virtualization Analysis	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
UNIX Hardware Variance	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Linux Hardware Variance	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
UNIX Shared OS (Zone) Analysis	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Wintel Hardware Variance	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Windows Shared OS (Zone) Analysis	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•

82

Consolidation & Virtualization Analysis Strategy Instrumentation Enablement

Key:

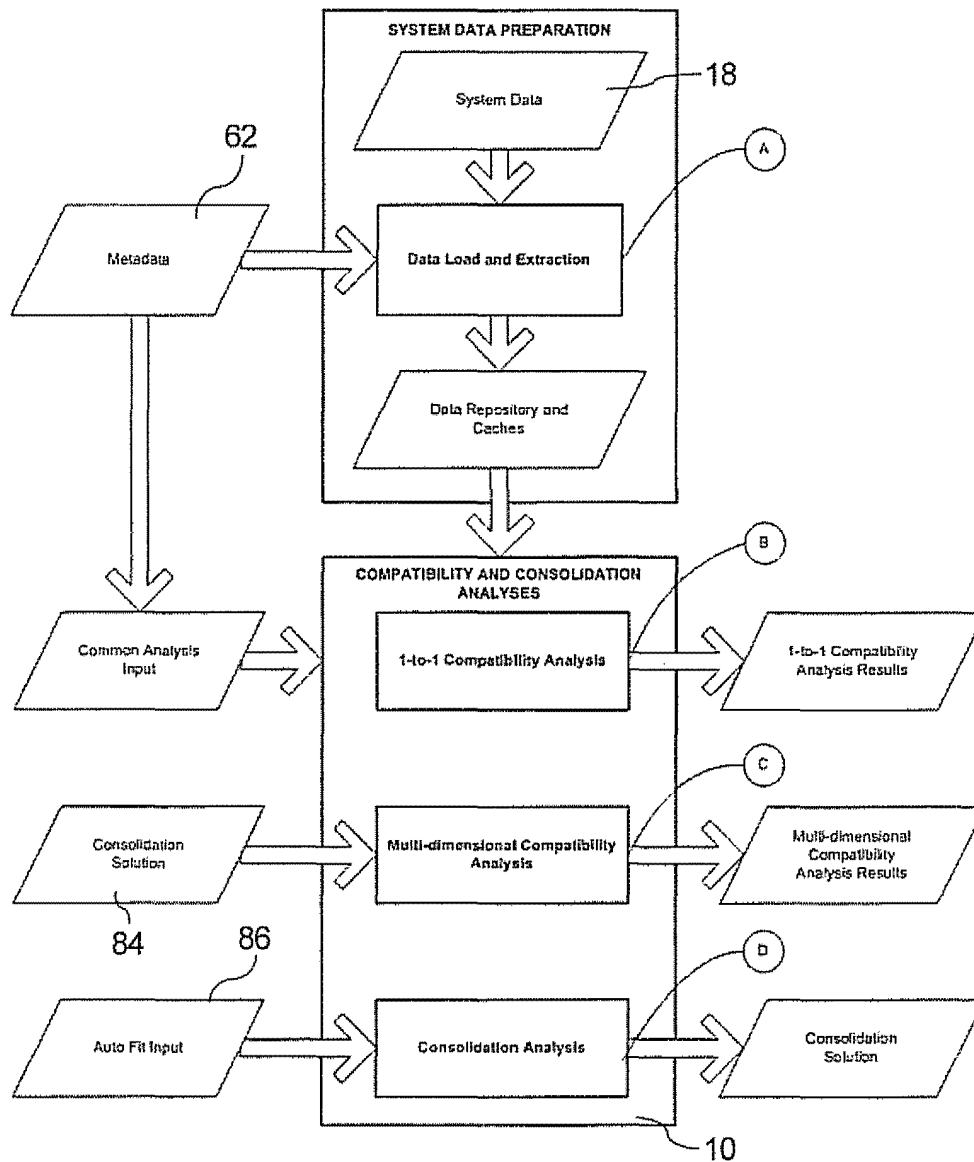
- Detailed Coverage
- Partial Coverage or Requires Customization

78

80

76

Figure 8



**Figure 9**

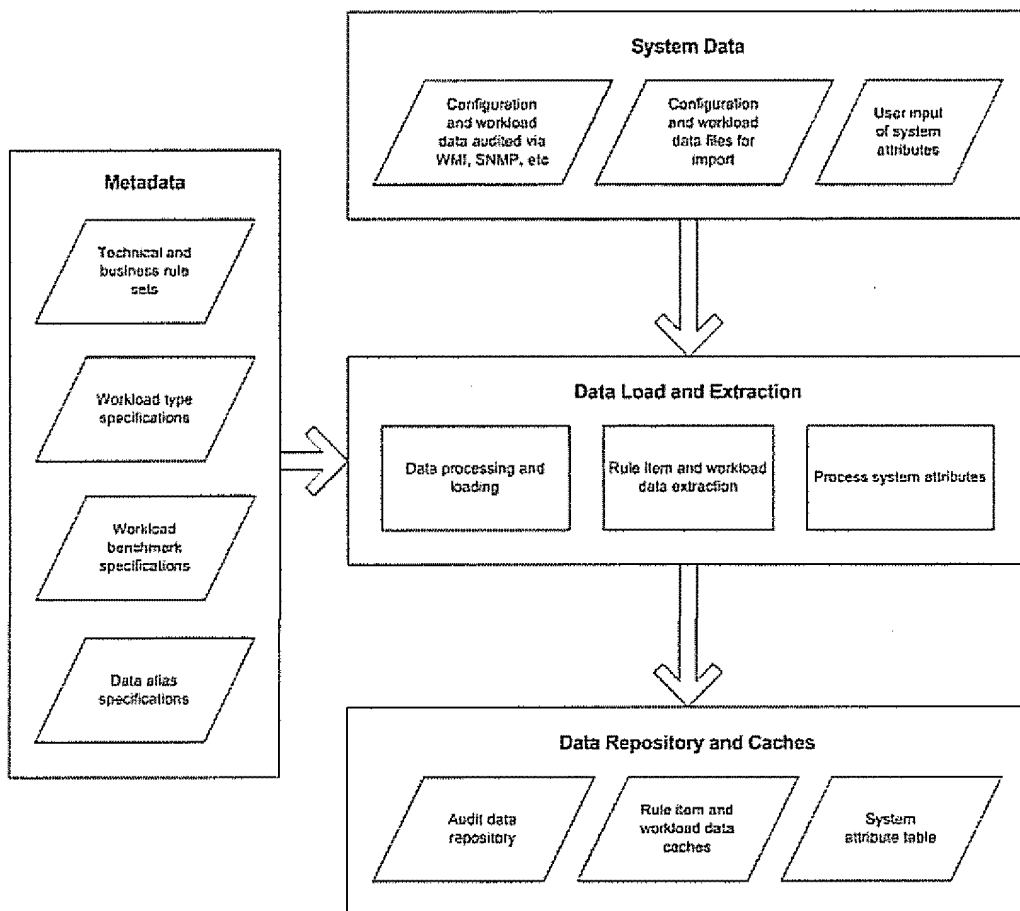


Figure 10

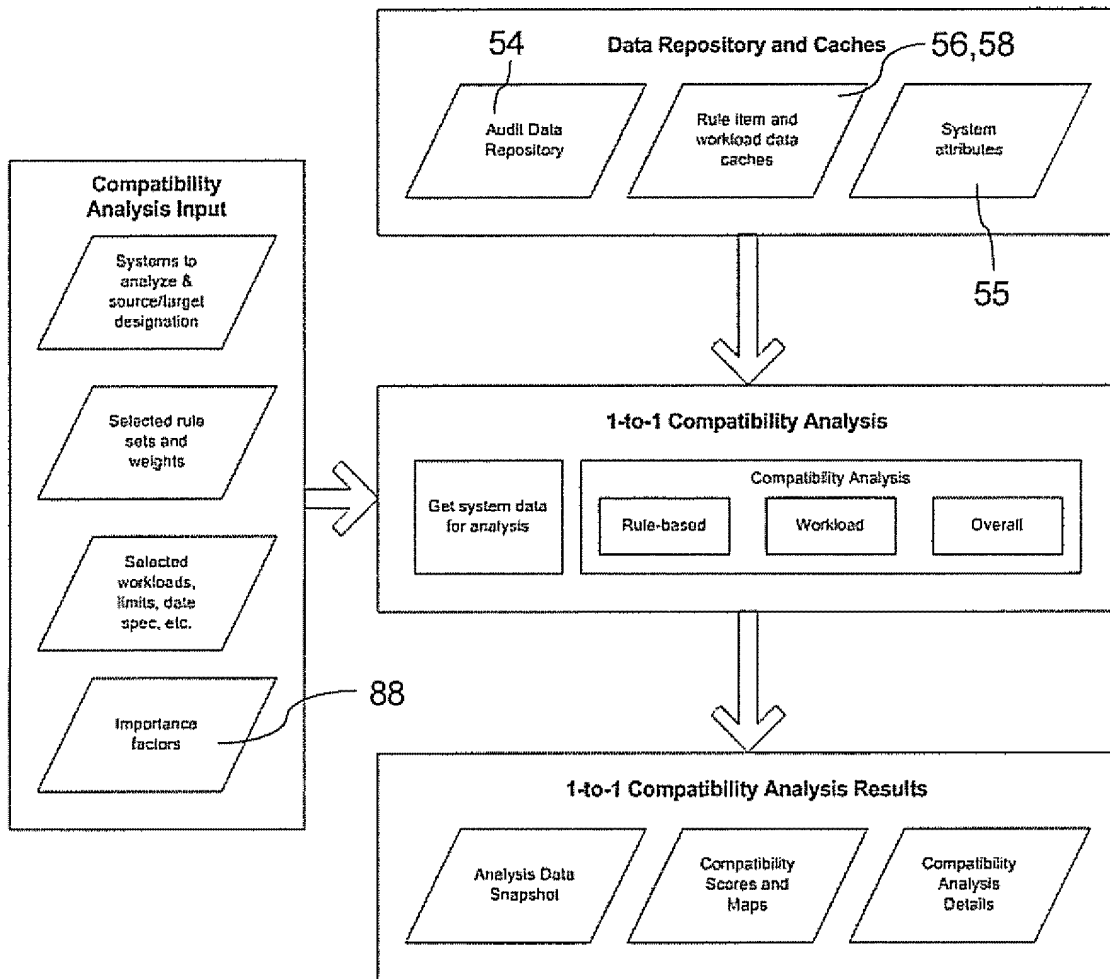


Figure 11

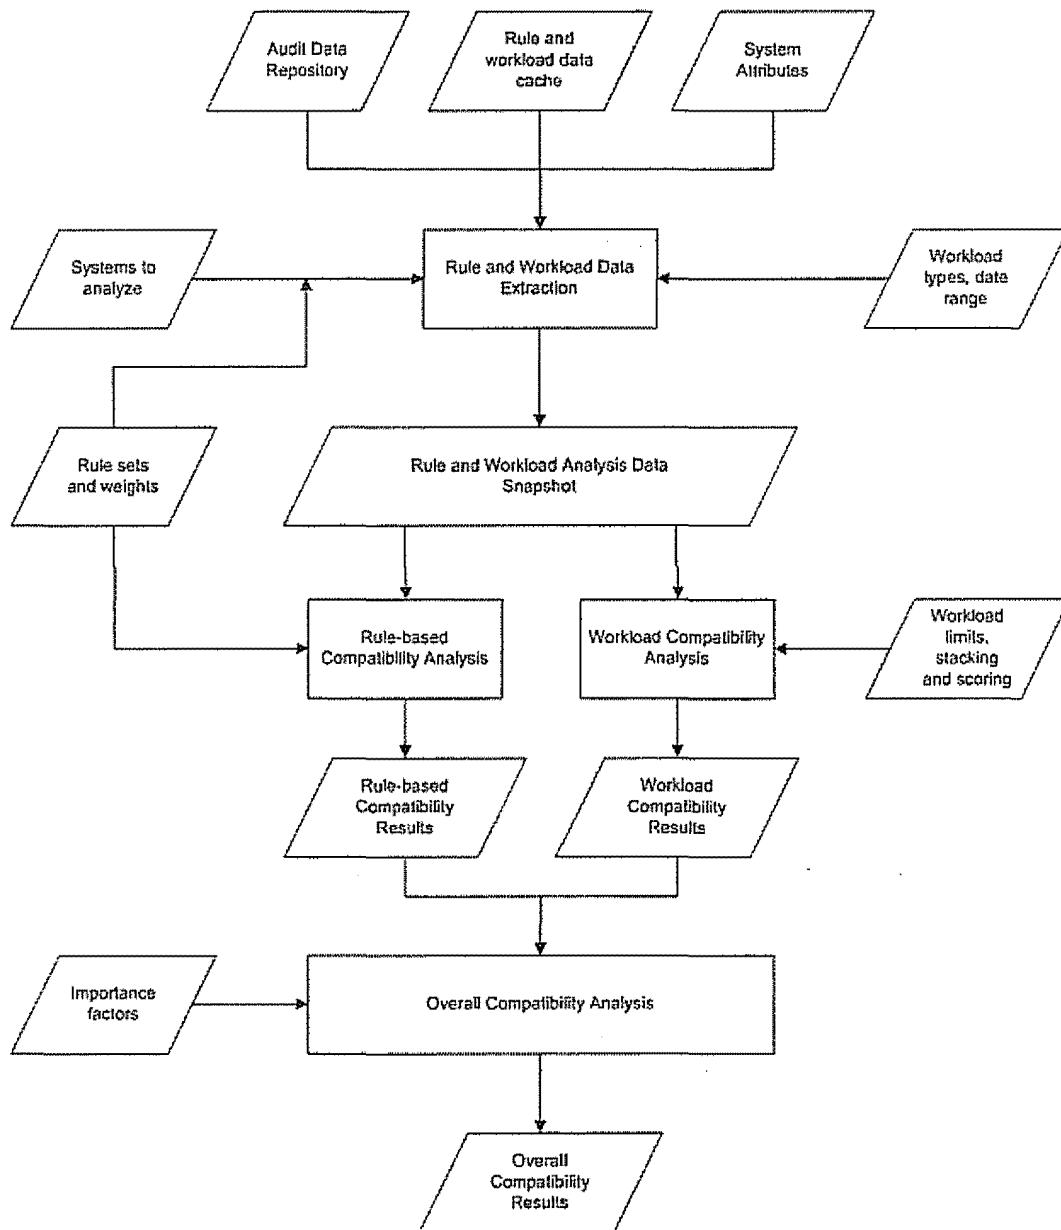


Name	Description
Solaris Kernel Settings	Detailed Analysis of Solaris Kernel Settings
Solaris Name Service Settings	Detailed Analysis of Solaris Name Service Settings
Solaris NVRAM Settings	Detailed Analysis of Solaris EEPROM Settings
UNIX Application Stacking	Detailed analysis rules for OS-level stacking of UNIX apps onto existing or new systems
UNIX Forward Consolidation	Detailed analysis rules for forward consolidation of UNIX applications onto existing infrastructure
WebLogic J2EE App Stacking	WebLogic J2EE Application Stacking Ruleset
Windows Compatibility	Assess compatibility between Windows systems
Acquisition Date	Group servers based on acquisition date
Application Tier	Prevent consolidating servers across application tiers.
Availability Target	Prevent combining servers with differing availability (uptime) targets
Backup Window	Prevent combining servers with differing backup windows
Building	Prevent consolidating servers between buildings
Business Service	Prevent consolidating servers across business services
Combined Constraints	Combined business constraint ruleset
Department	Prevent consolidating servers across departments
DR Strategy	Prevent combining servers with inconsistent DR strategies
Lease Renewal Date	Group servers based on lease renewal date
Location	Prevent consolidating servers across physical locations
Maintenance Window	Prevent consolidating servers with contentious maintenance windows
Operational Environment	Prevent consolidating servers between environments
Operational Role	Prevent consolidating servers that have differing functional roles
Ownership	Prevent combining servers that belong to different owners
Primary Application	Prevent combining servers that are servicing different applications
Service Plan	Prevent consolidating servers that have different service plans
UPS Profile	Prevent combining servers that have differing UPS requirements
Oracle Stacking Analysis Data Validation	Determine data coverage for app stacking analysis
UNIX Analysis Data Validation	Determine data coverage for app stacking analysis
Windows Analysis Data Validation	Determine data coverage for windows stacking and virtualization analysis
Oracle Data Stacking	Assess data consolidation potential between Oracle instances
Oracle Instance Stacking	Assess affinity for instance-level stacking
SQL Server 2000 (WMI)	Map out data-level consolidation opportunities in MS SQL 2000 environments
Windows File and Print Consolidation	Analyze compatibility of file and print servers
Hardware Models	Map out manufacturers and models
Macro Analysis	Combined Macro Analysis ruleset
OS Map	Map out OS Types, Versions and kernel patch levels
OS Map (with Kernel Bits)	Map out high-level OS affinity regions

**Figure 12**

Name	Description
CPU Utilization	CPU resource constraints
Disk I/O	Disk I/O constraints
Network I/O	Network I/O constraints
Disk capacity	Available disk space
Memory	Available memory resources

Figure 13



**Figure 14**

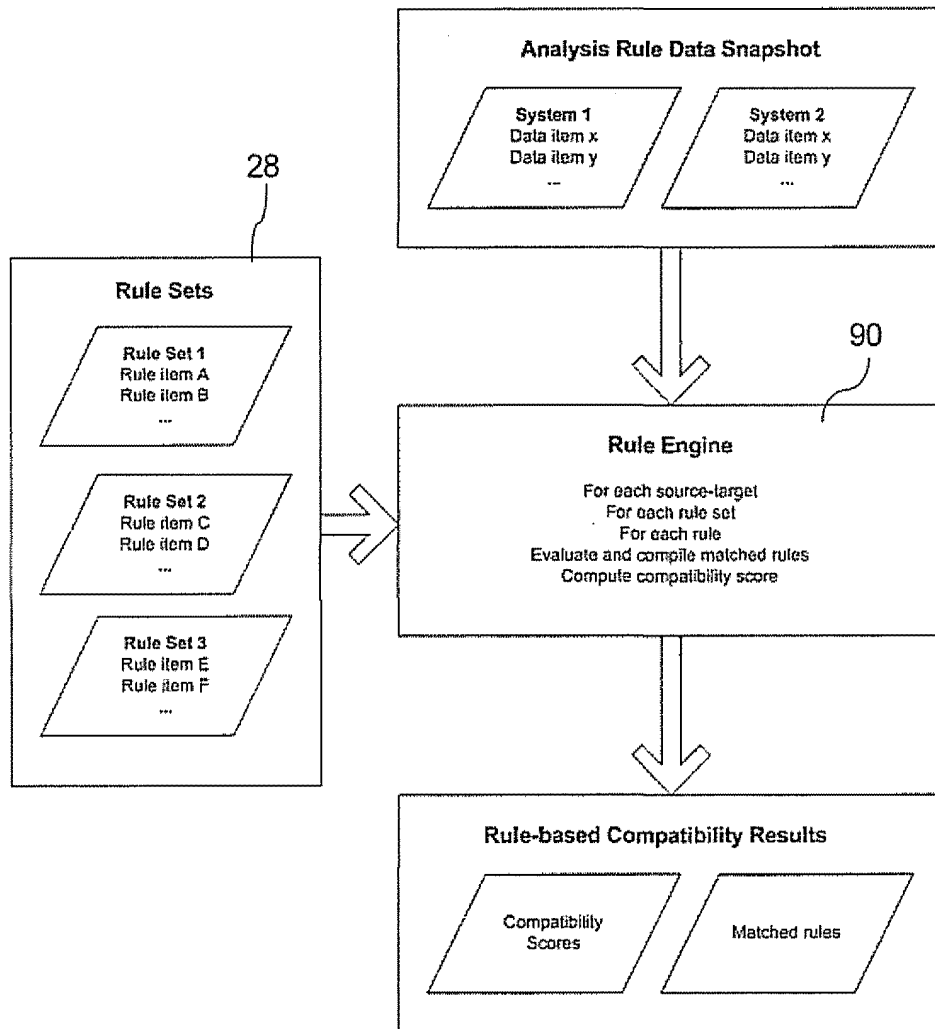


Figure 15

	Rule Type	Rule Specifier	Source	Target	Weight	Mutex Flag	Match Flag	Suppress Flags	Remediation Cost	Description
1	AliasQuery	OS			70%		OS		\$1,200	Different Operating Systems
2	AliasQuery	OS Version			30%		VER	OS	\$320	Different Operating System Versions
3	AliasQuery	Time Zone			20%				\$320	Systems are in Different Time Zones
4	AliasQuery	Total memory		<	5%				\$350	Target Has Less Memory
5	UriQuery	cirba-cm/ServerInformation/OSPatchLevel			5%		PATC H	OS VER	\$160	Different Patch Levels
6	UriQuery	cirba-pa/patchesTable#*			5%	Y		OS VER	\$80	Patch Differences Between Systems
7	UriQuery	cirba-cm/ServerInformation/OSKernelBits			10%			OS	\$80	Not Running Same Kernel Bits
8	UriQuery	cirba-cm/KernelParametersTable#*	present	absent	5%	Y		OS	\$80	Some Kernel Parameters are Absent
9	UriQuery	cirba-cm/KernelParametersTable/Setting#SHMMAX		<	5%				\$80	Maximum Shared Memory Settings Differ
10	UriQuery	cirba-si/appInvTableTable.appInvTableTable/version#oracle	9	8	4%		ORA		\$1,500	Different Version of Oracle
11	UriQuery	cirba-si/appInvTableTable.appInvTableTable/version#oracle			5%			ORA	\$3,500	Different Version of Oracle
12	UriQuery	cirba-si/appInvTableTable.appInvTableTable/version#apache			2%		PATC H		\$80	Different Version of Apache
13	UriQuery	win32-os/Win32_OperatingSystem/ServicePackMajorVersion			5%			OS VER	\$160	Different Service Pack Levels
14	UriQuery	win32-quickfix/win-hotfix#*	present	absent	5%	Y		OS VER	\$80	Patch Differences Between Systems
15	UriQuery	win32-os/Win32_StartupCommand#*			2%	Y		OS	\$80	Startup Commands are Different
16	UriQuery	win32-misc/Win32_Environment/VariableValue#Path			2%			OS	\$80	Global Path is Different
17	UriQuery	win32-misc/Win32_Environment/VariableValue#Path <SYSTEM>			2%			OS	\$80	System Path is Different
18	UriQuery	win32-os/Win32_Service#*	present	absent	5%	Y		OS	\$160	Some Services not On Target
19	UriQuery	win32-os/Win32_Service/Started			1%	Y		OS	\$80	Some Services not Started On Target
20	UriQuery	win32-application/win-products/McAfee_VirusScan_Enterprise	present	absent	5%			OS	\$200	McAfee Not On Target

Figure 16

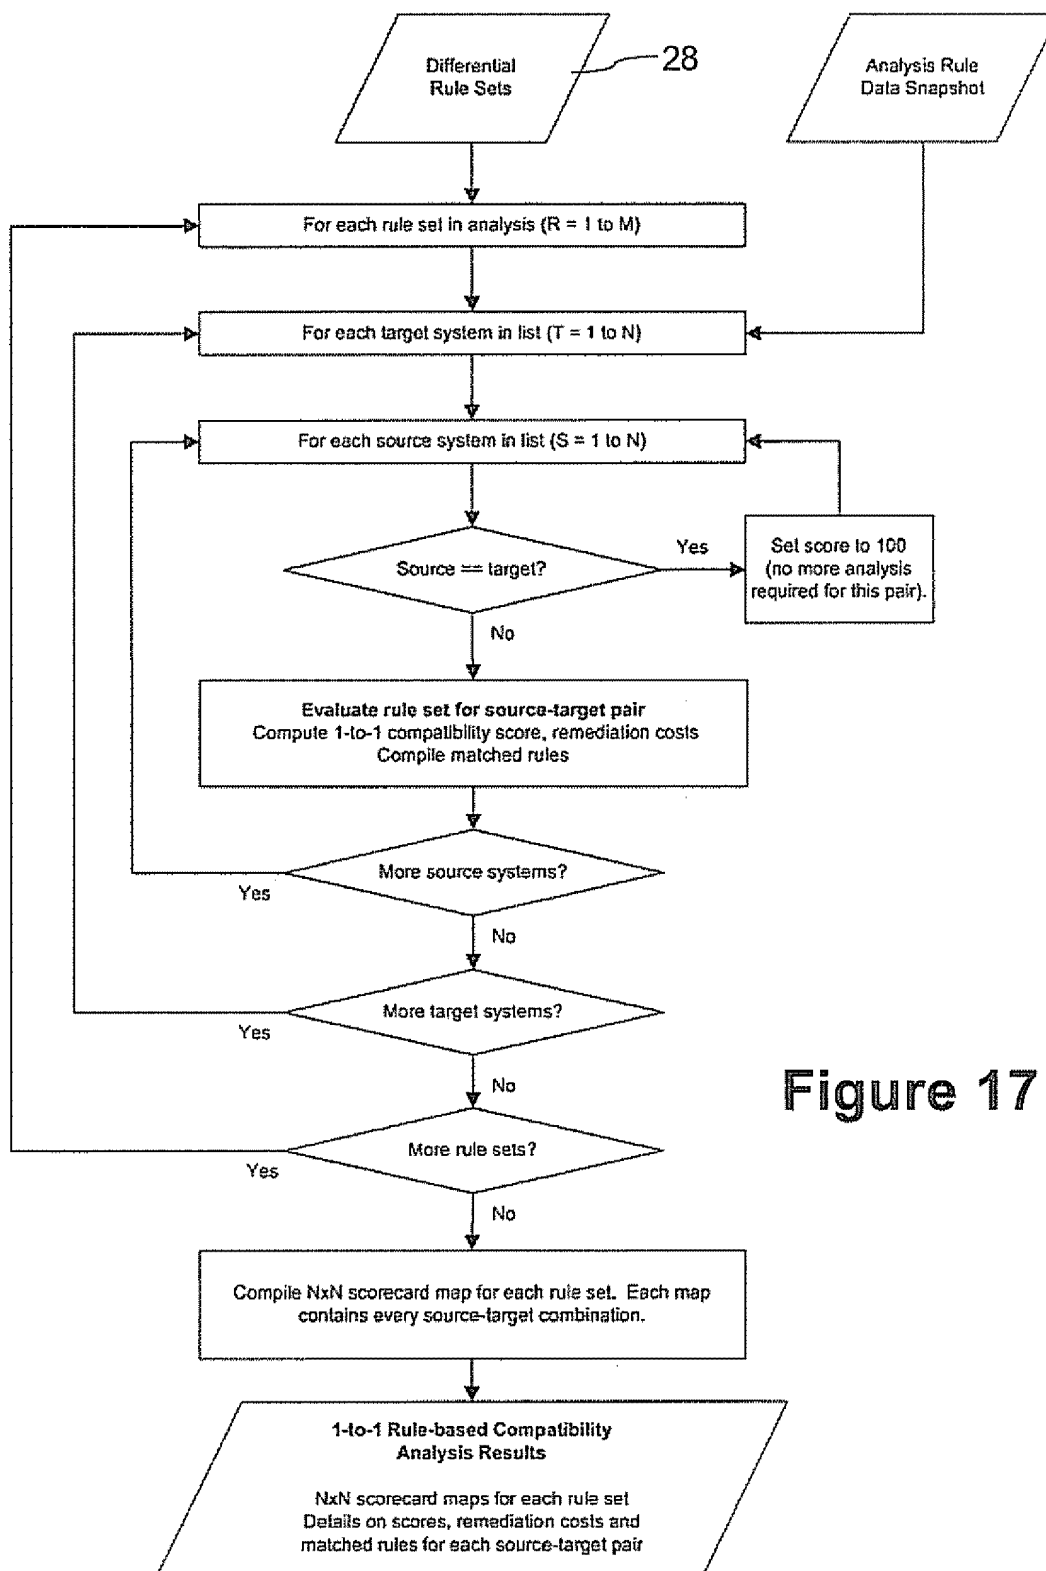
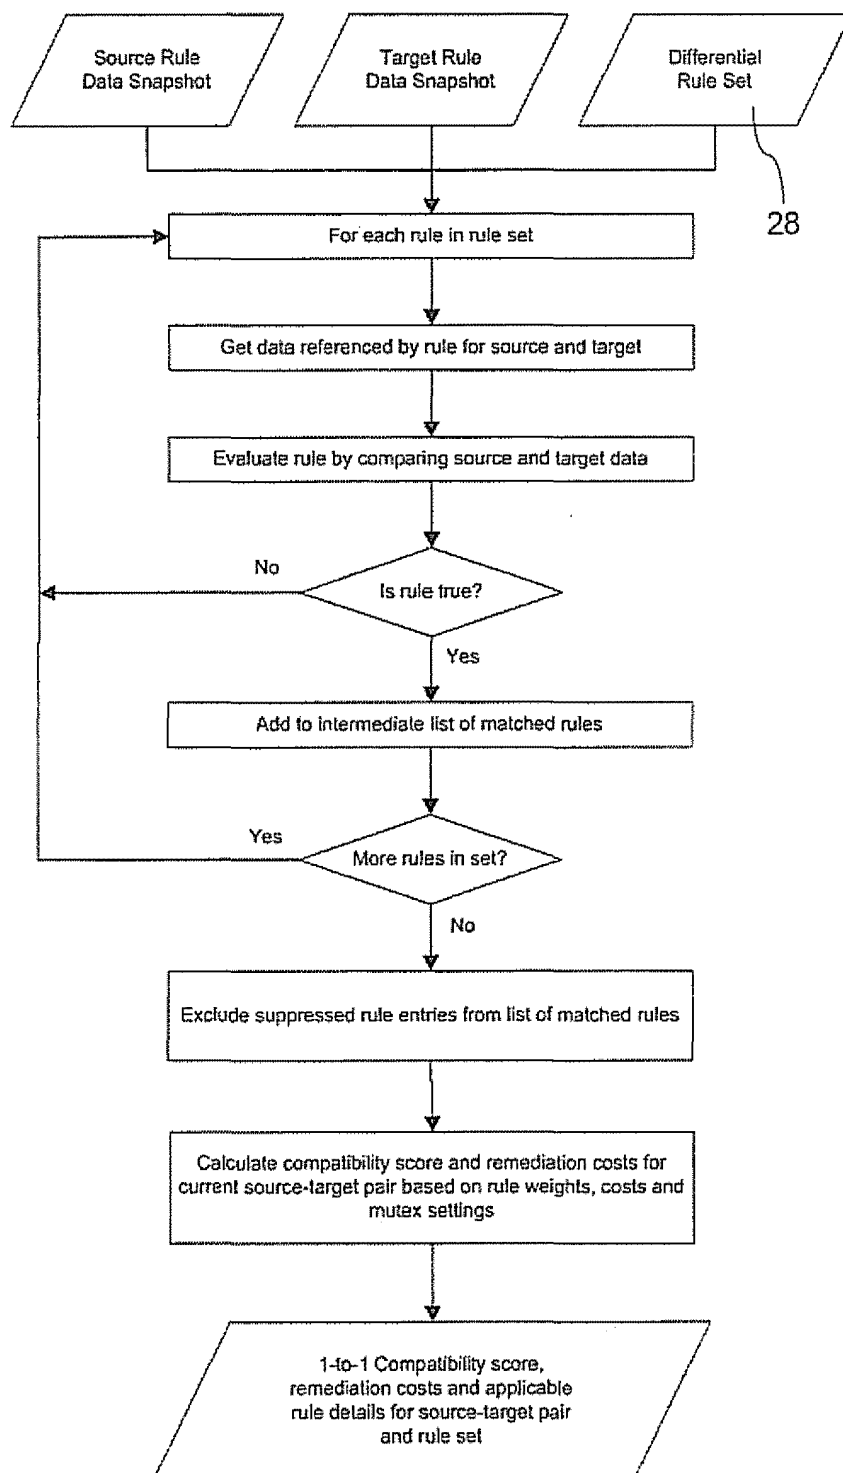


Figure 17



**Figure 18**

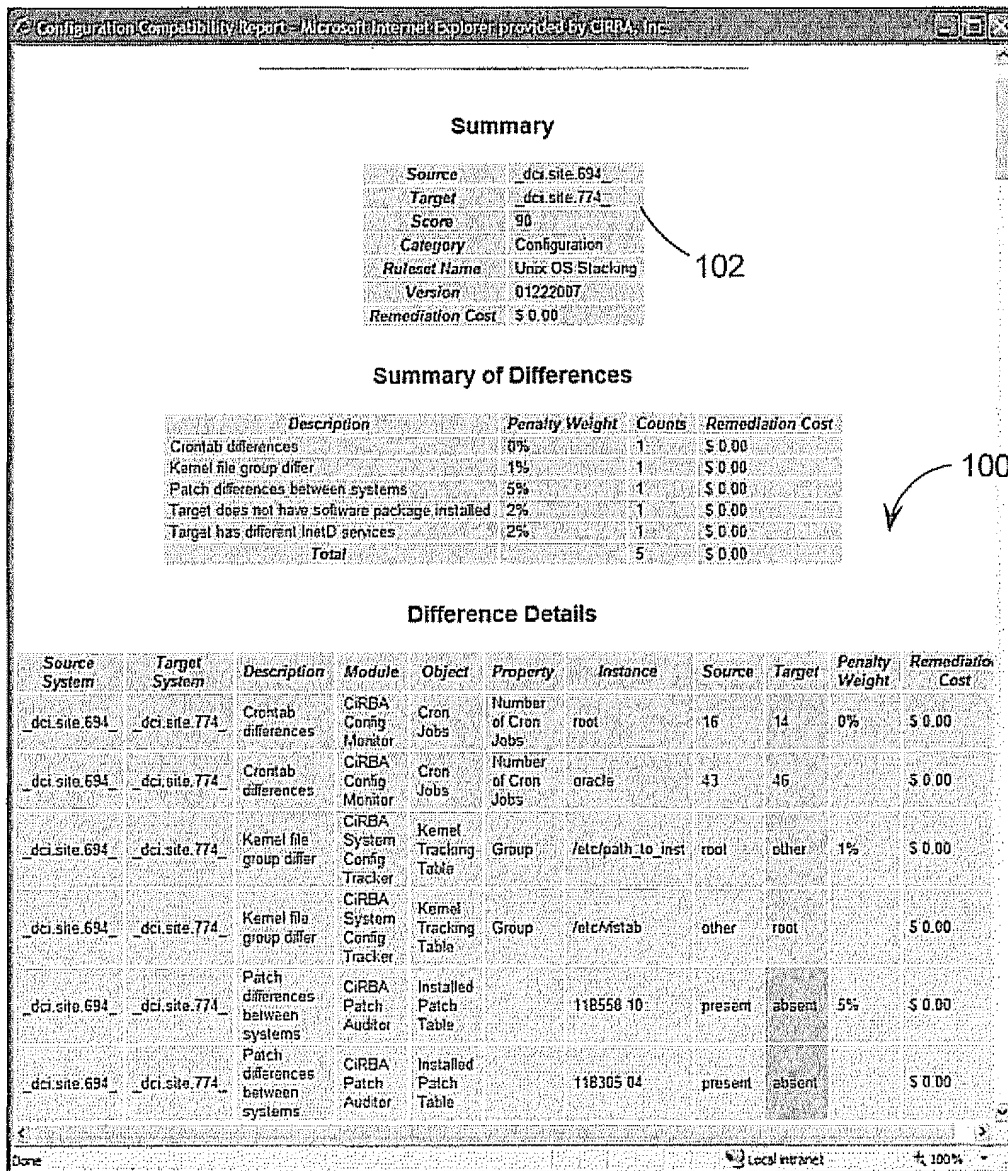
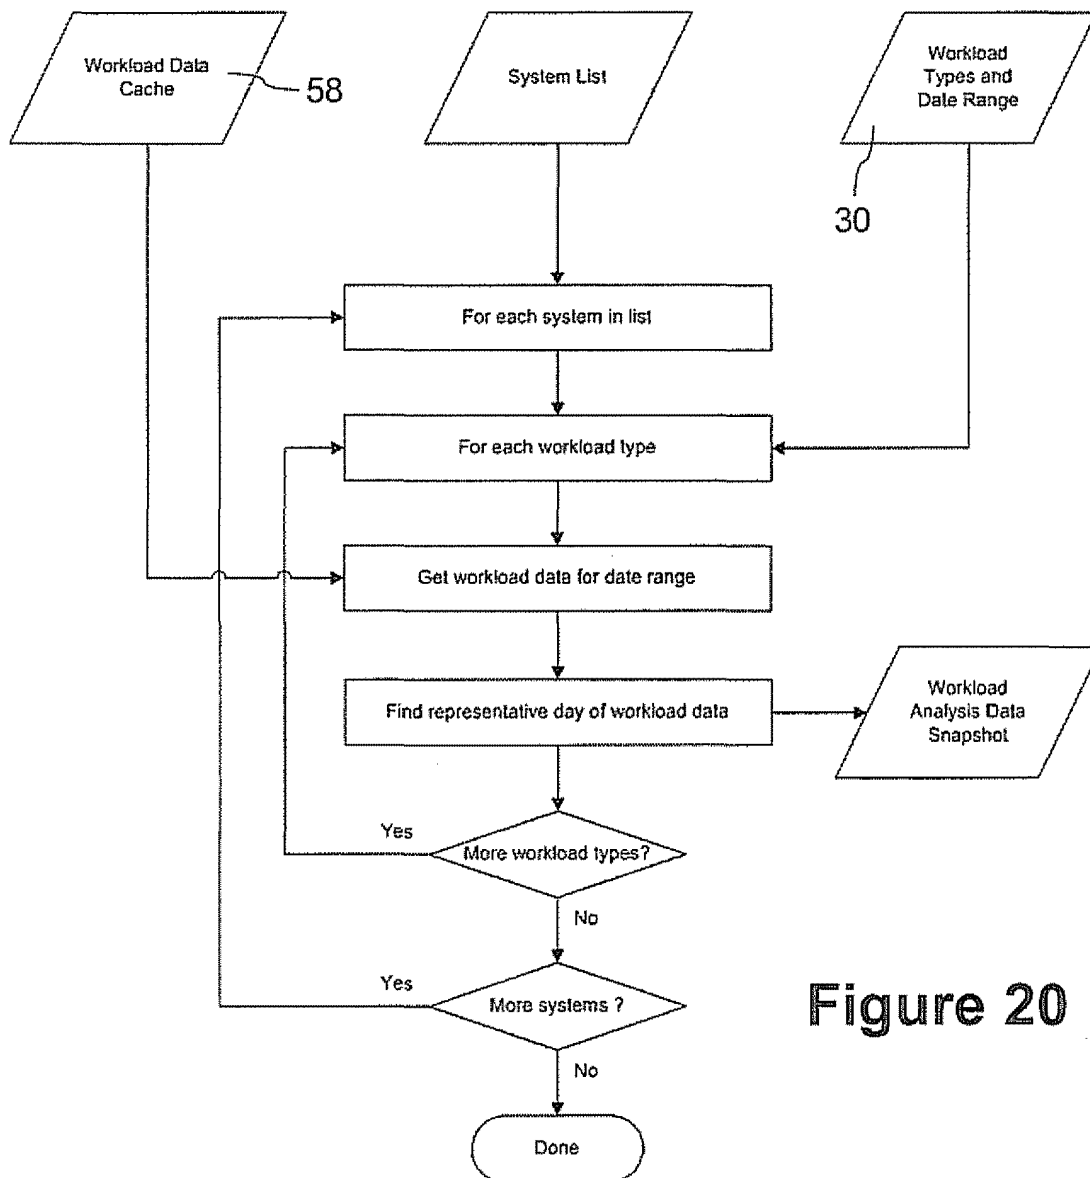
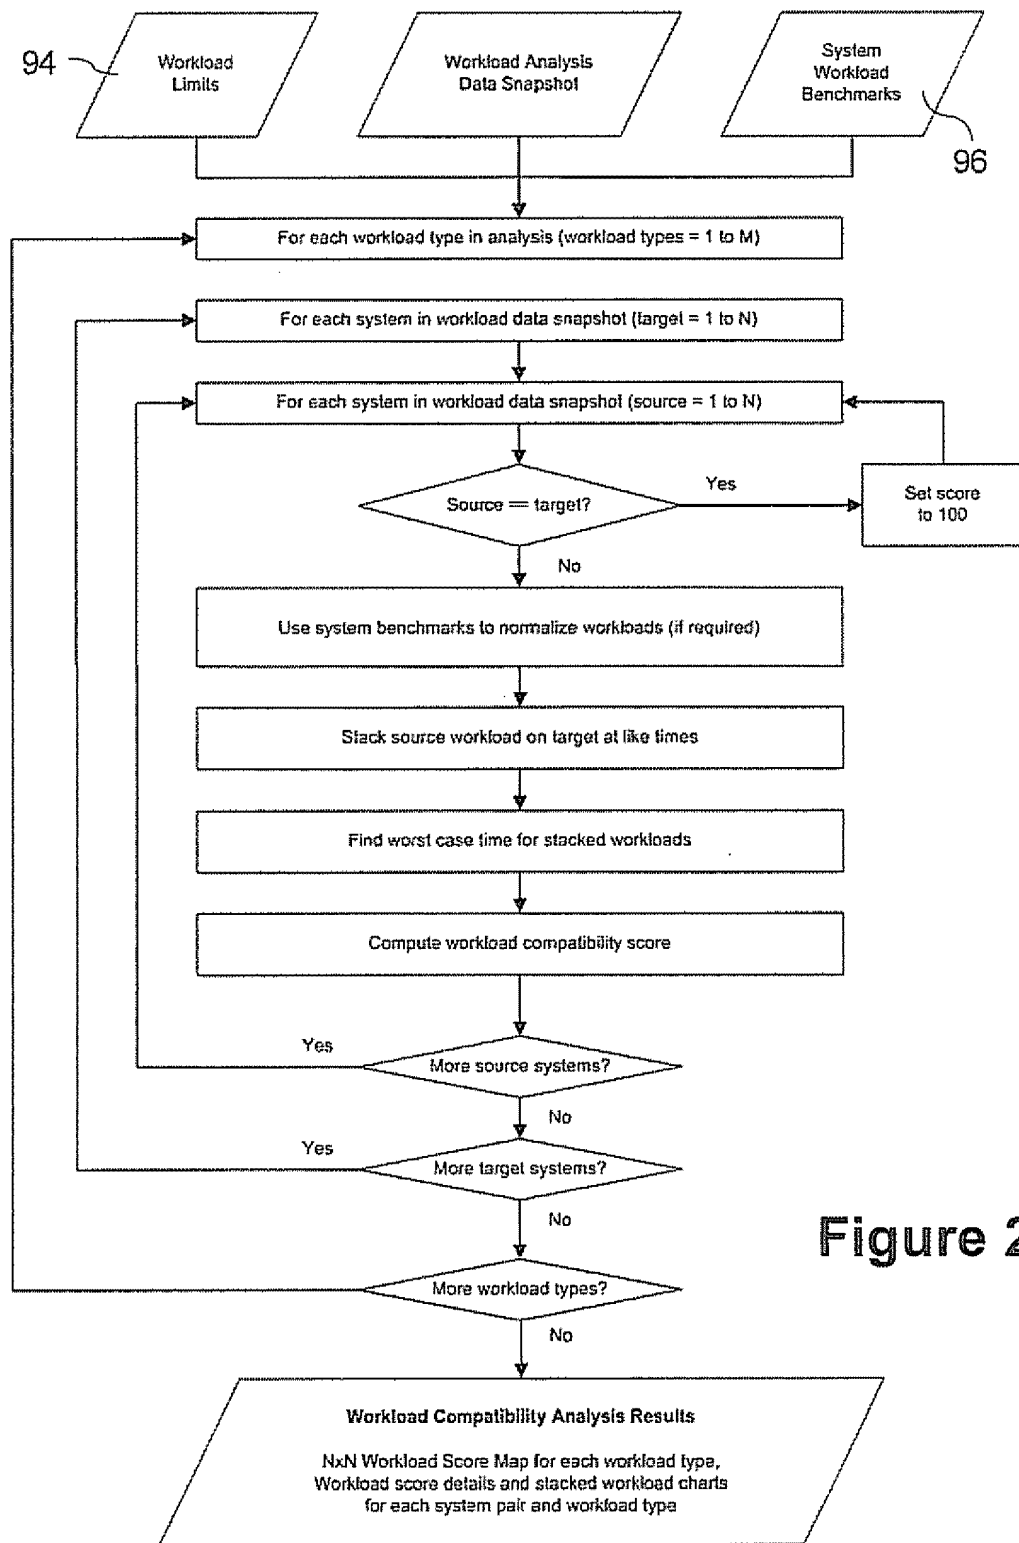


Figure 19





**Figure 20**



**Figure 21**


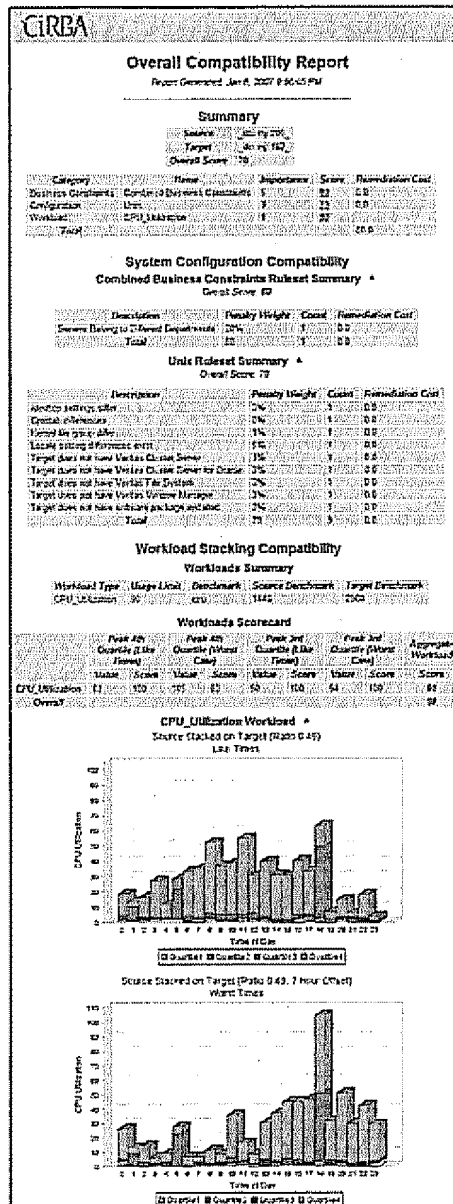


Figure 22



98

Figure 23

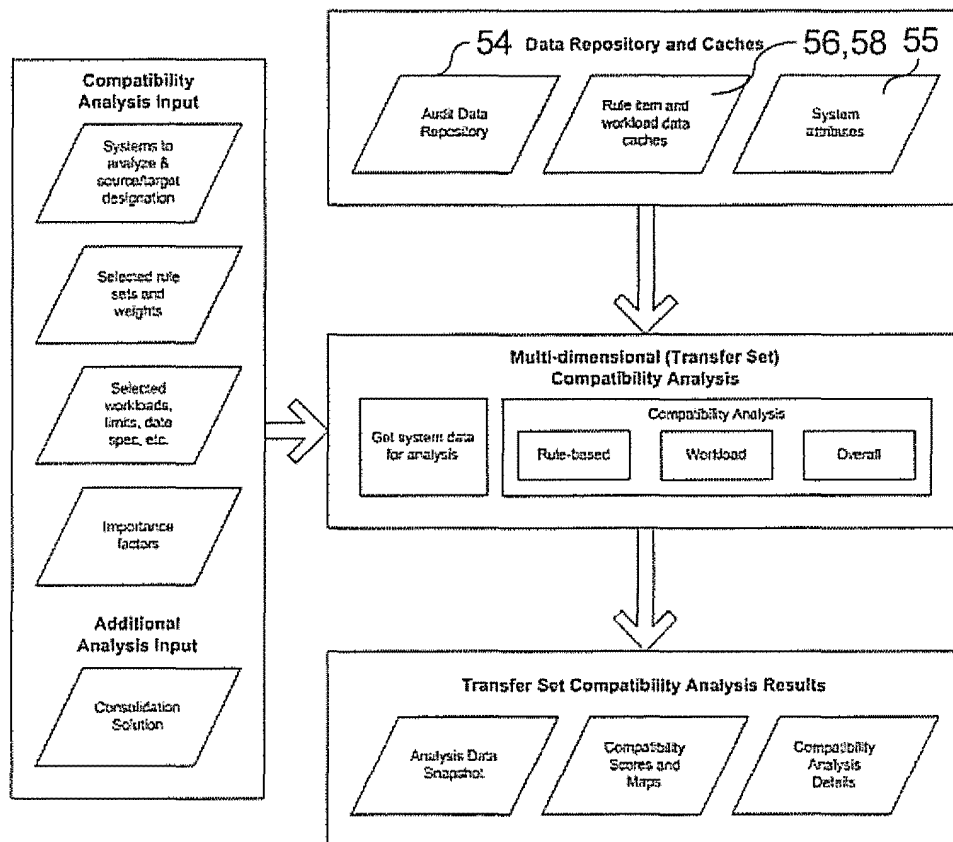
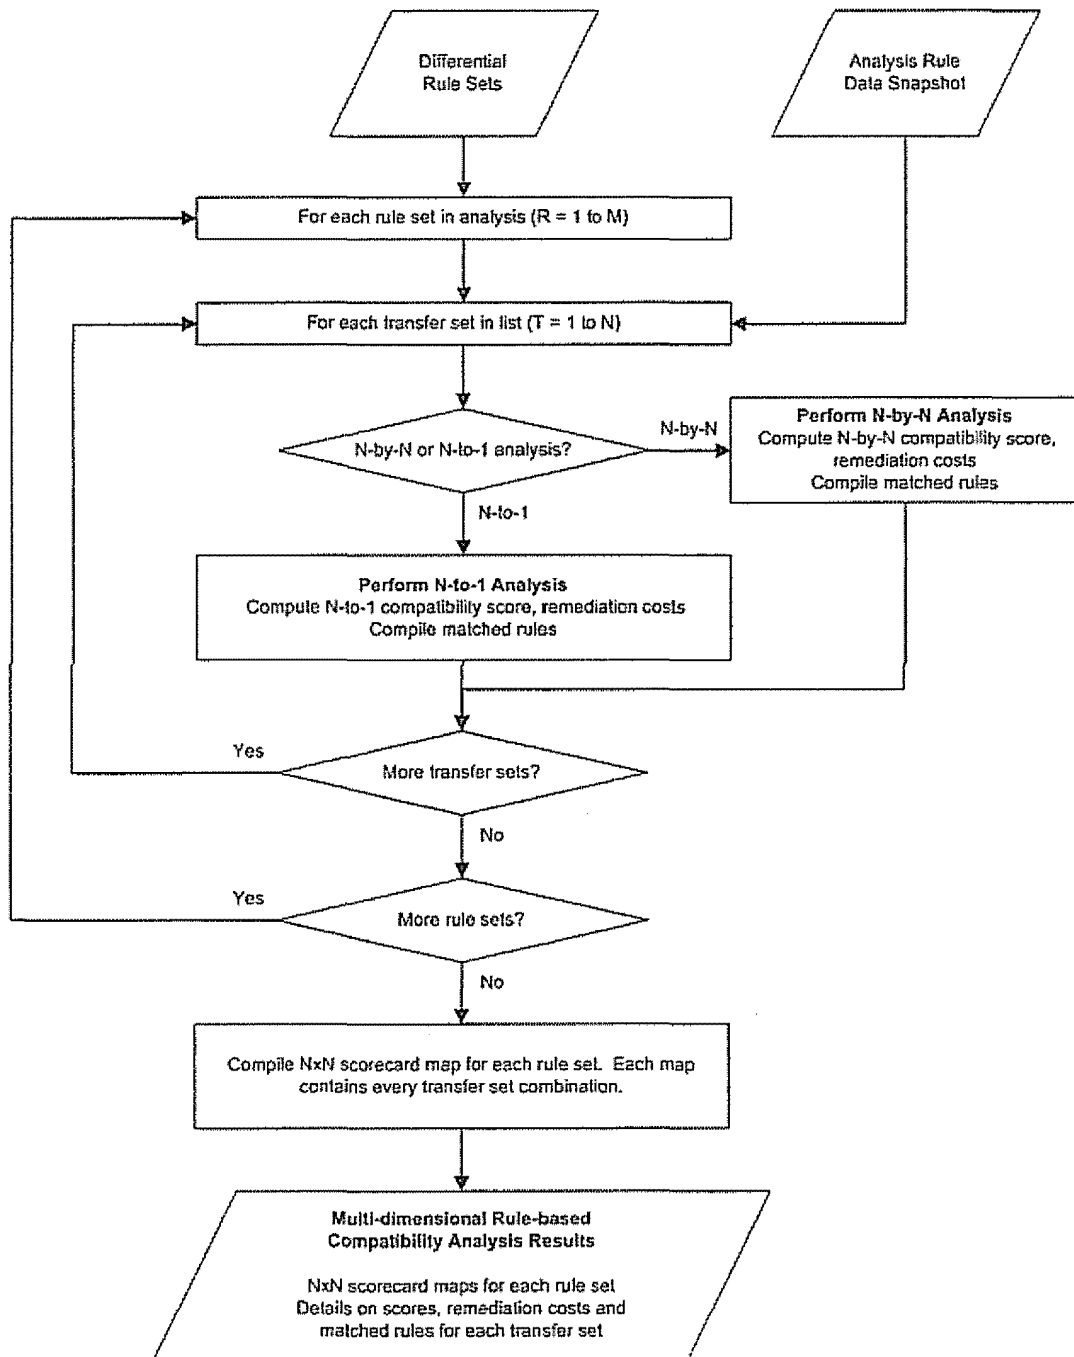
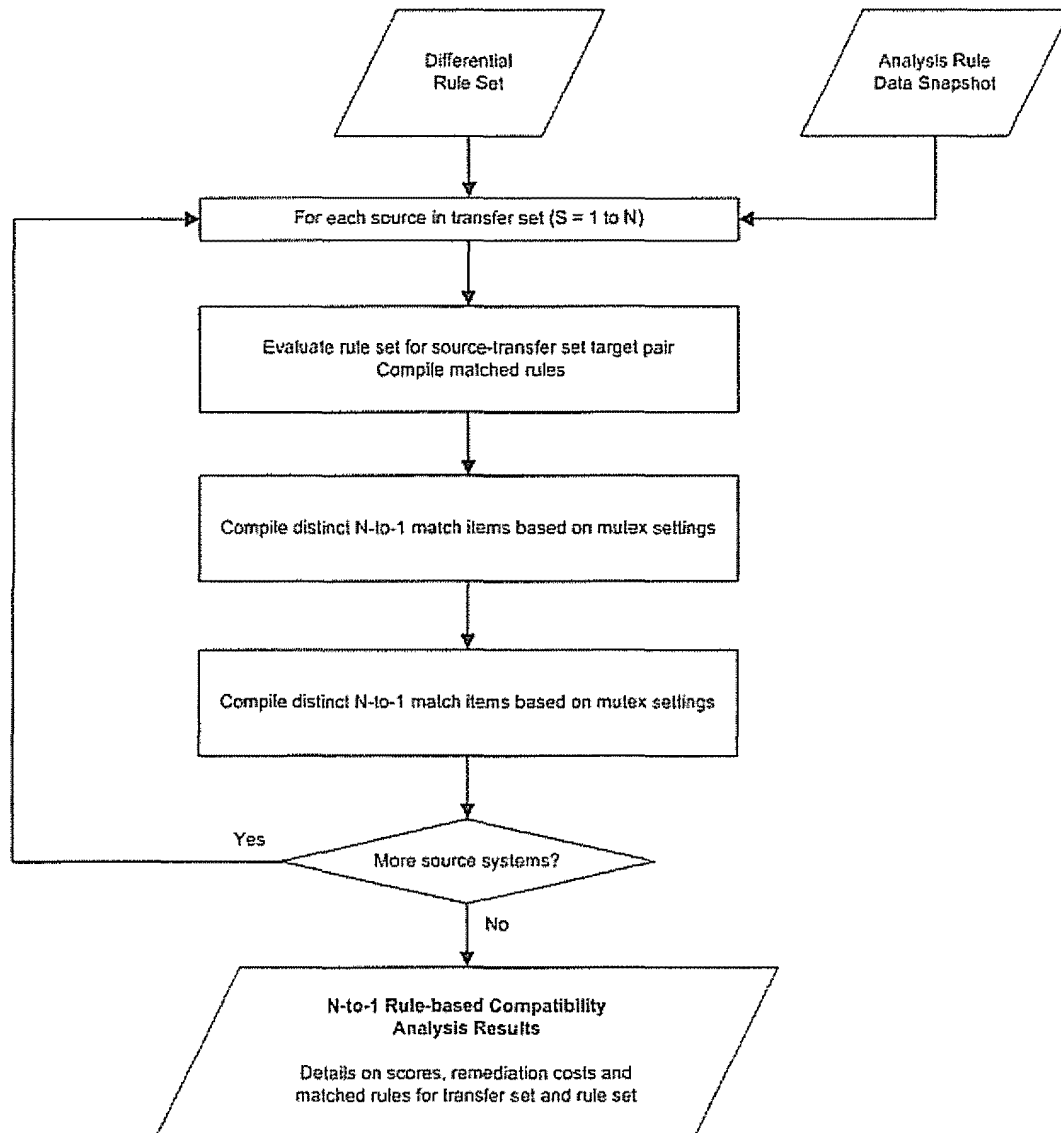


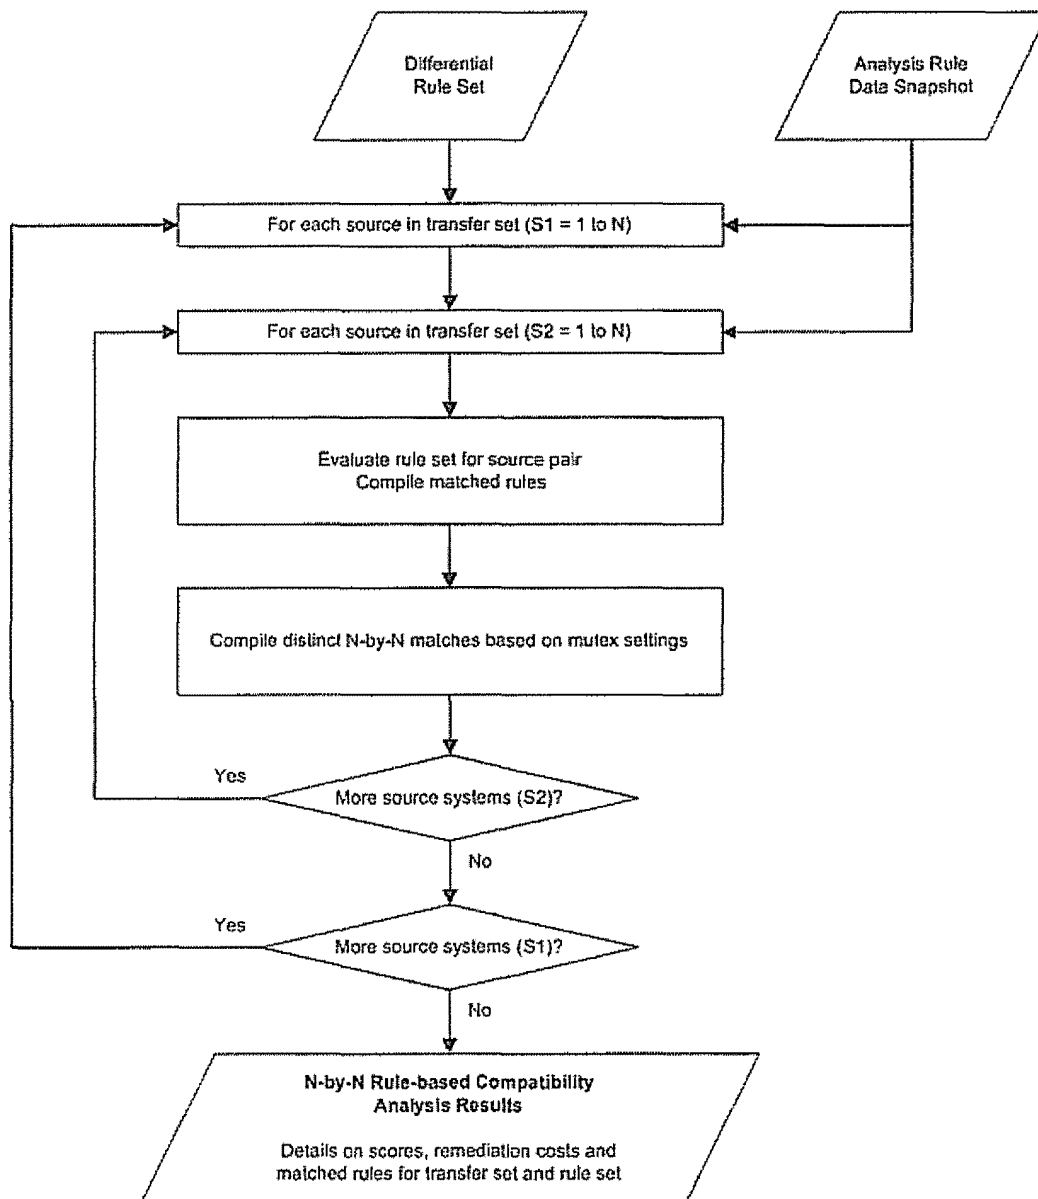
Figure 24(a)



**Figure 24(b)**

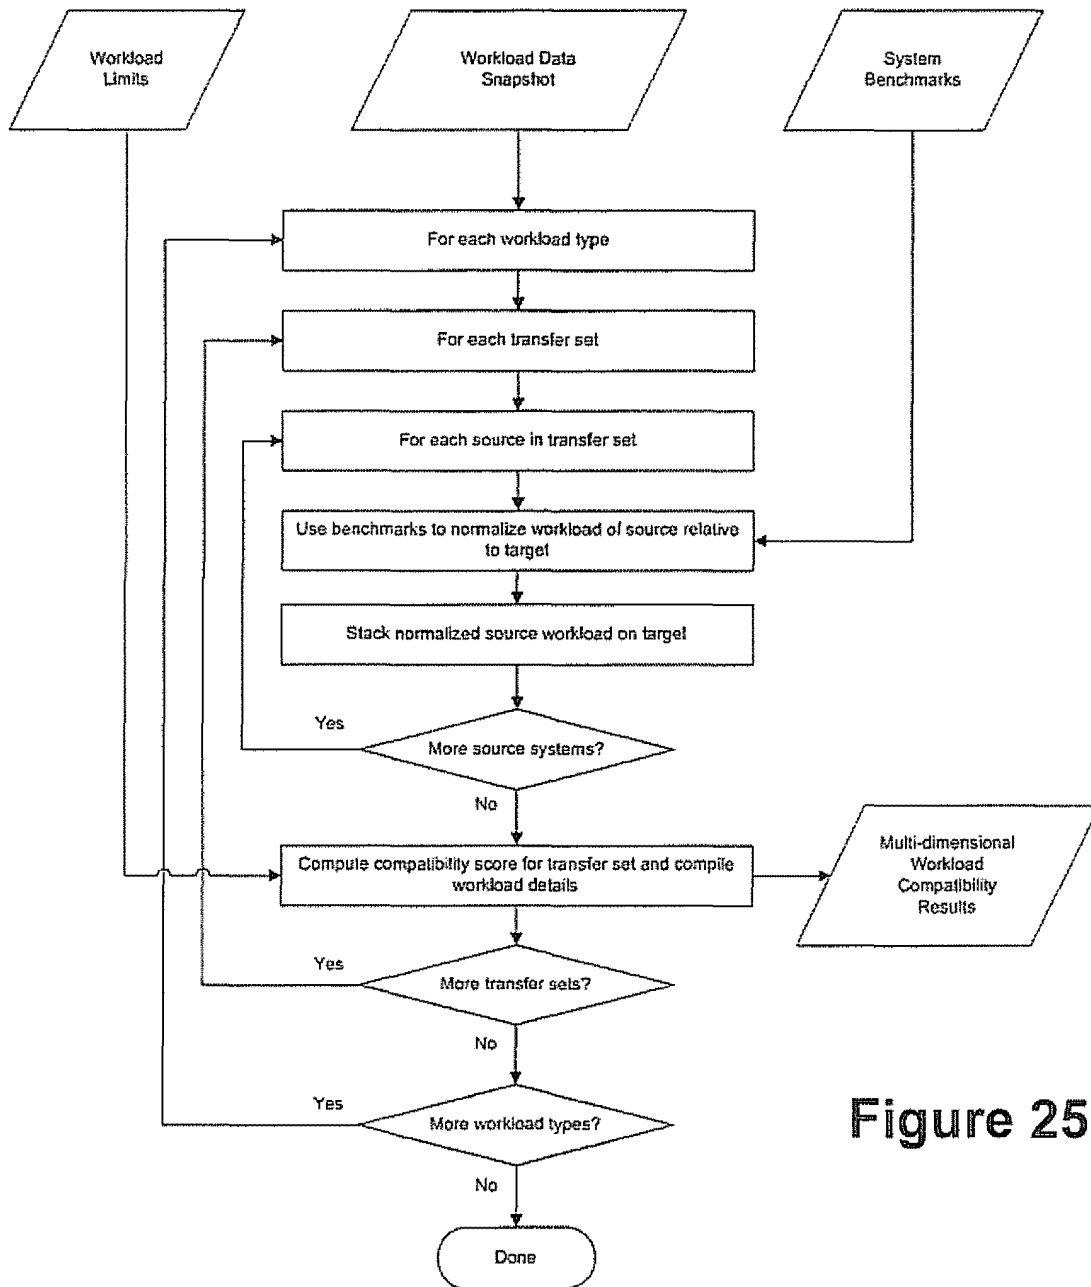


**Figure 24(c)**



**Figure 24(d)**





**Figure 25**



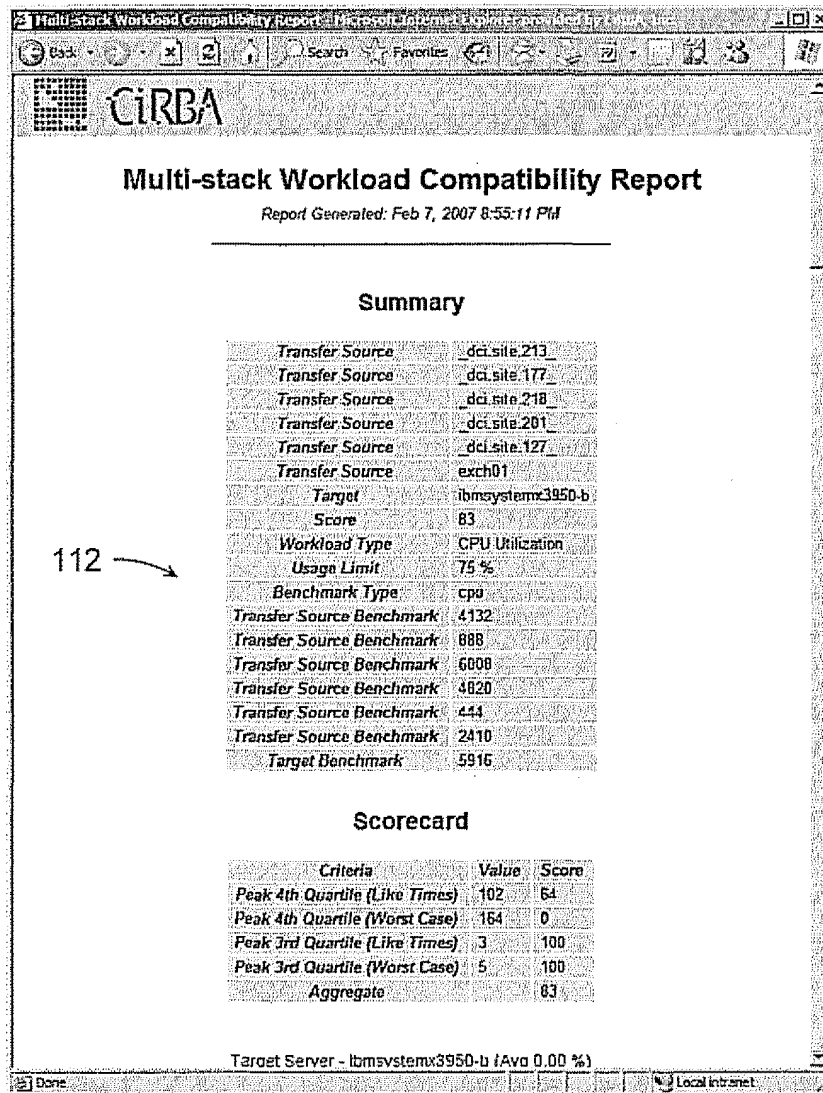


Figure 27




Figure 28

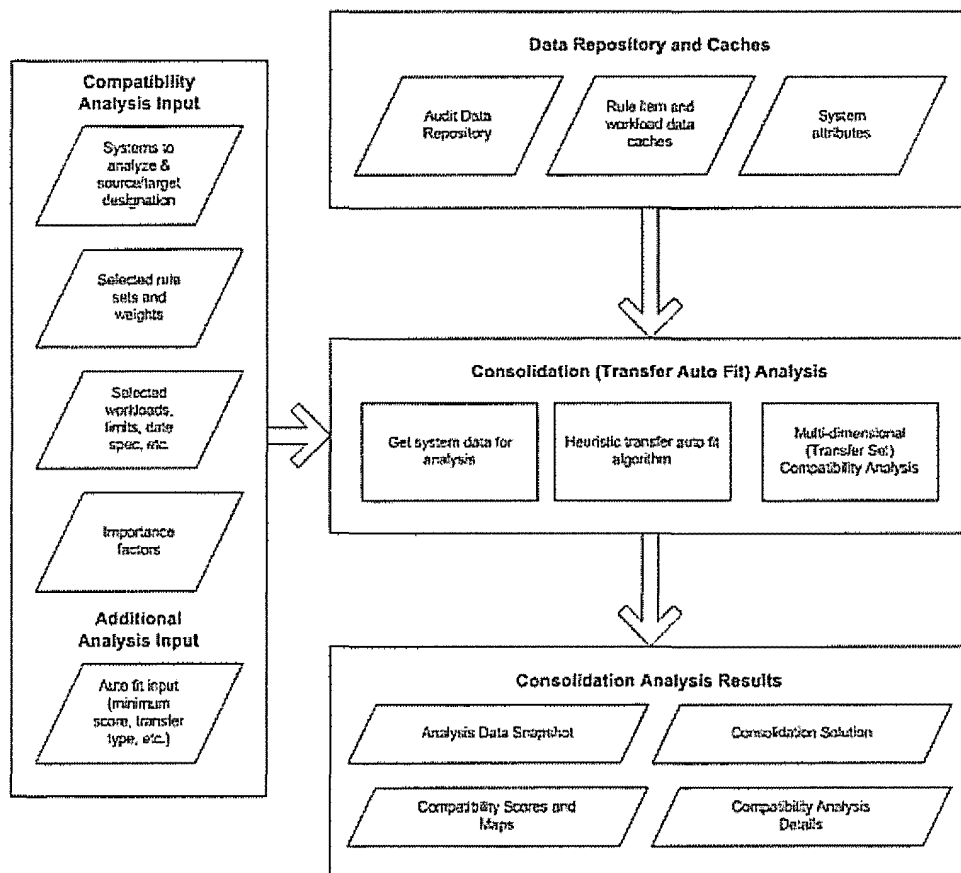


Figure 29

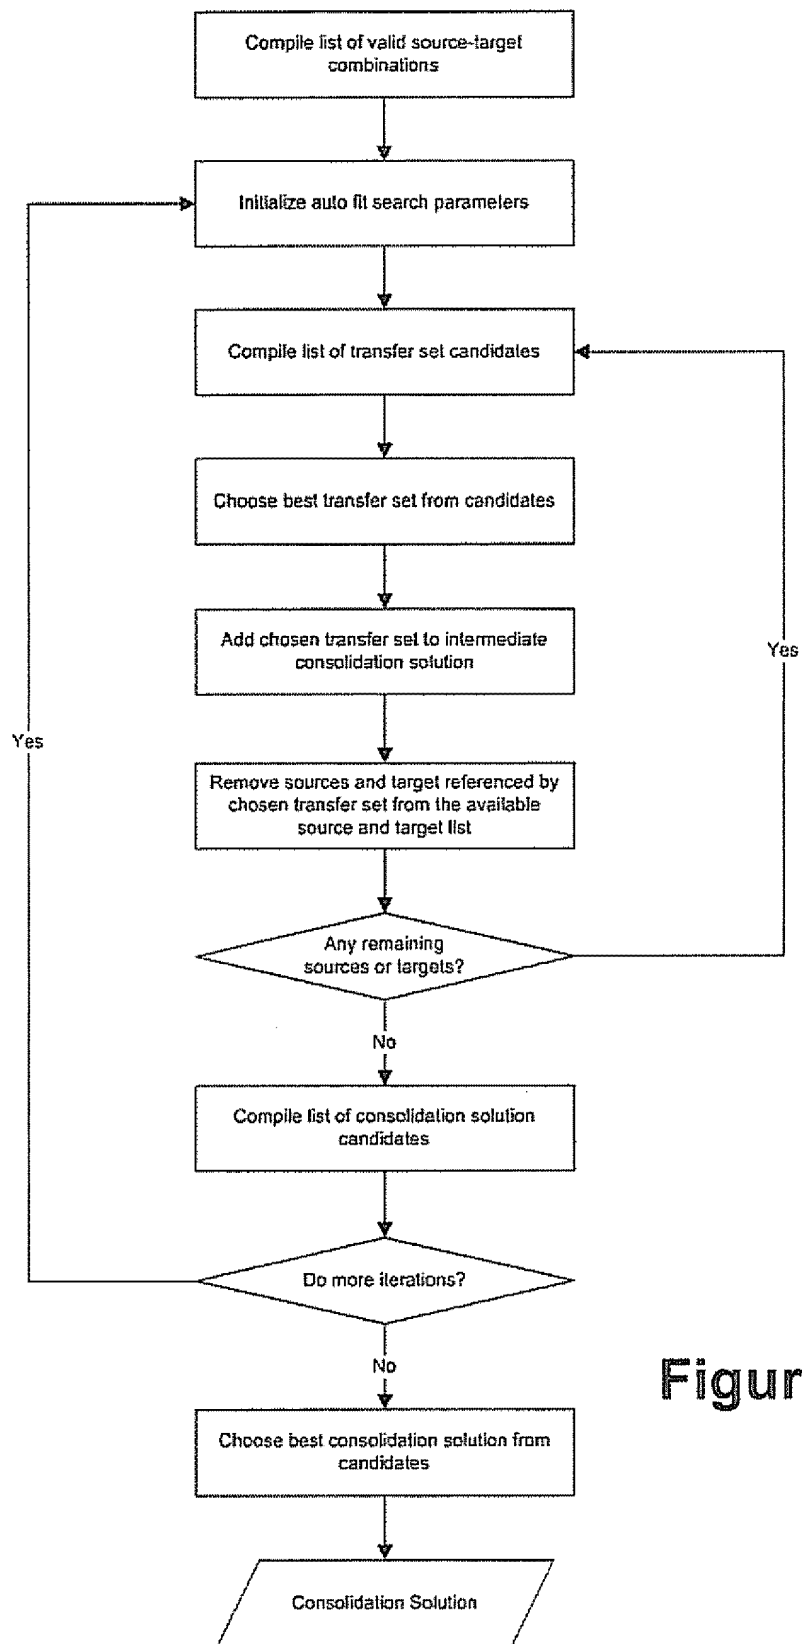
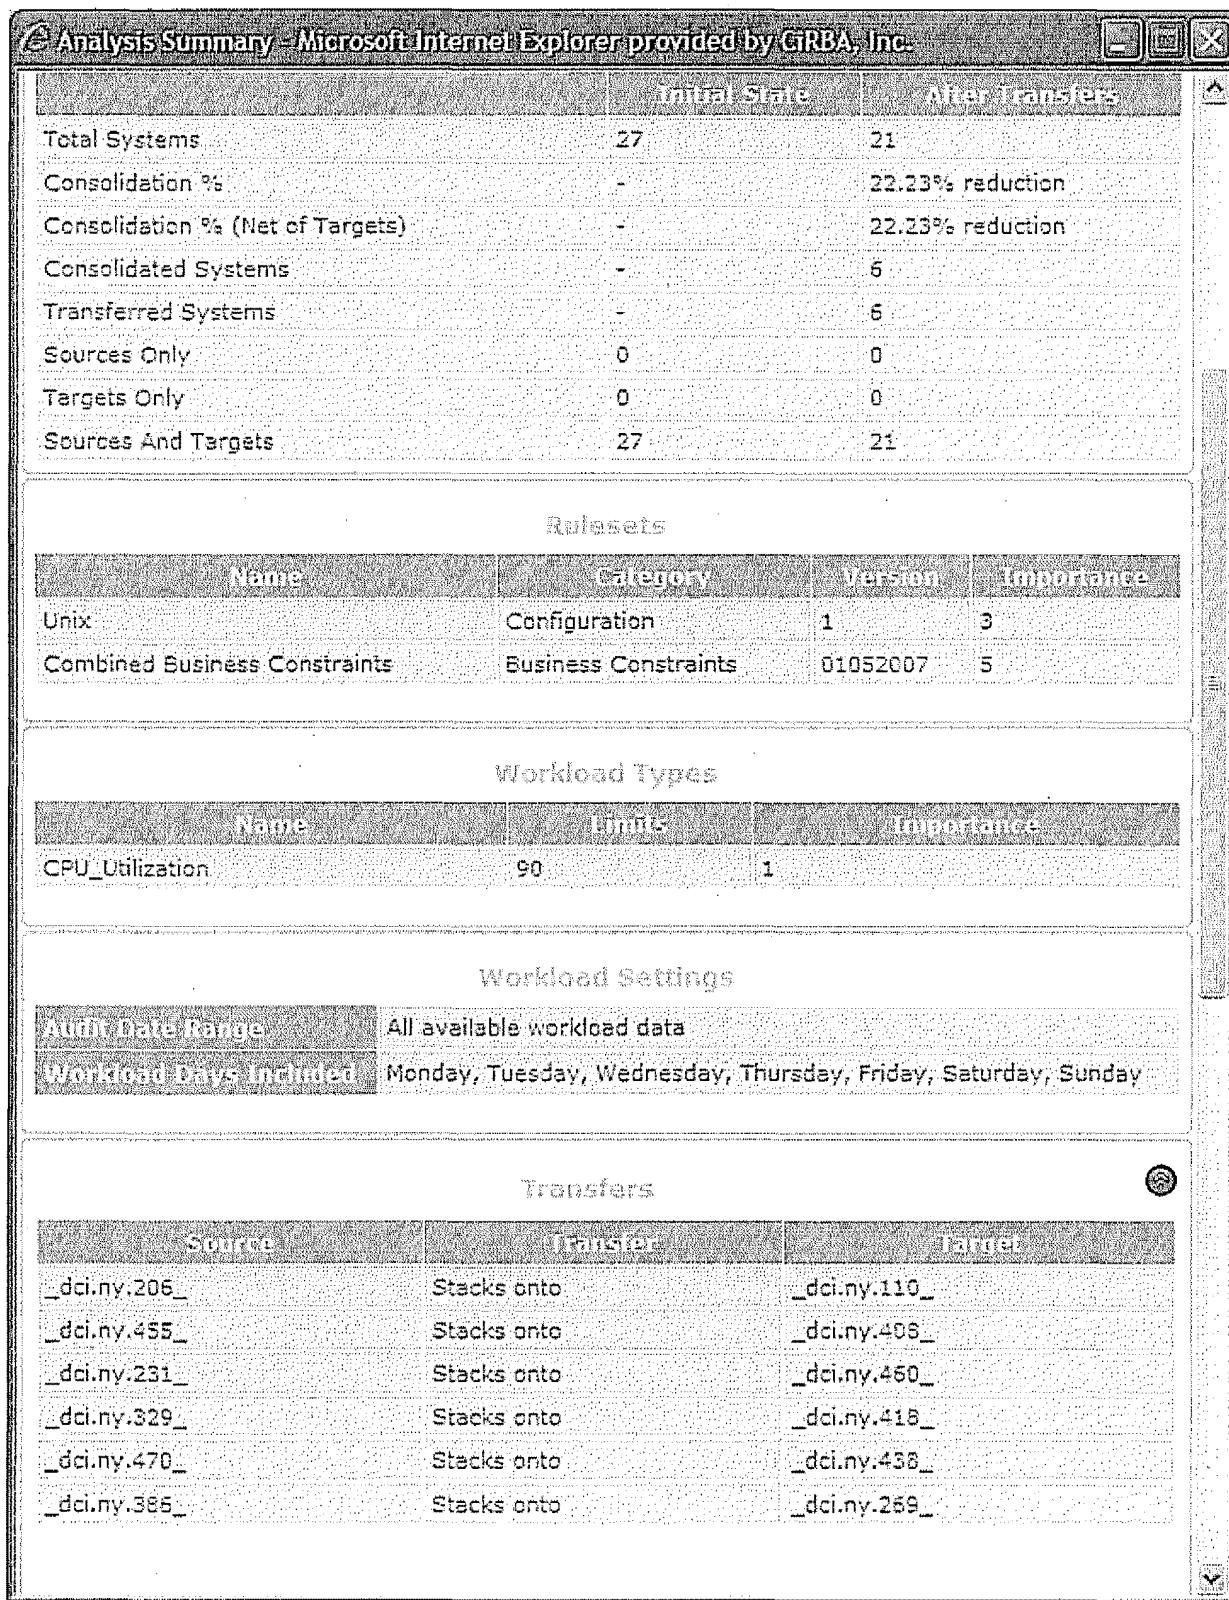


Figure 30



116 →

Figure 31

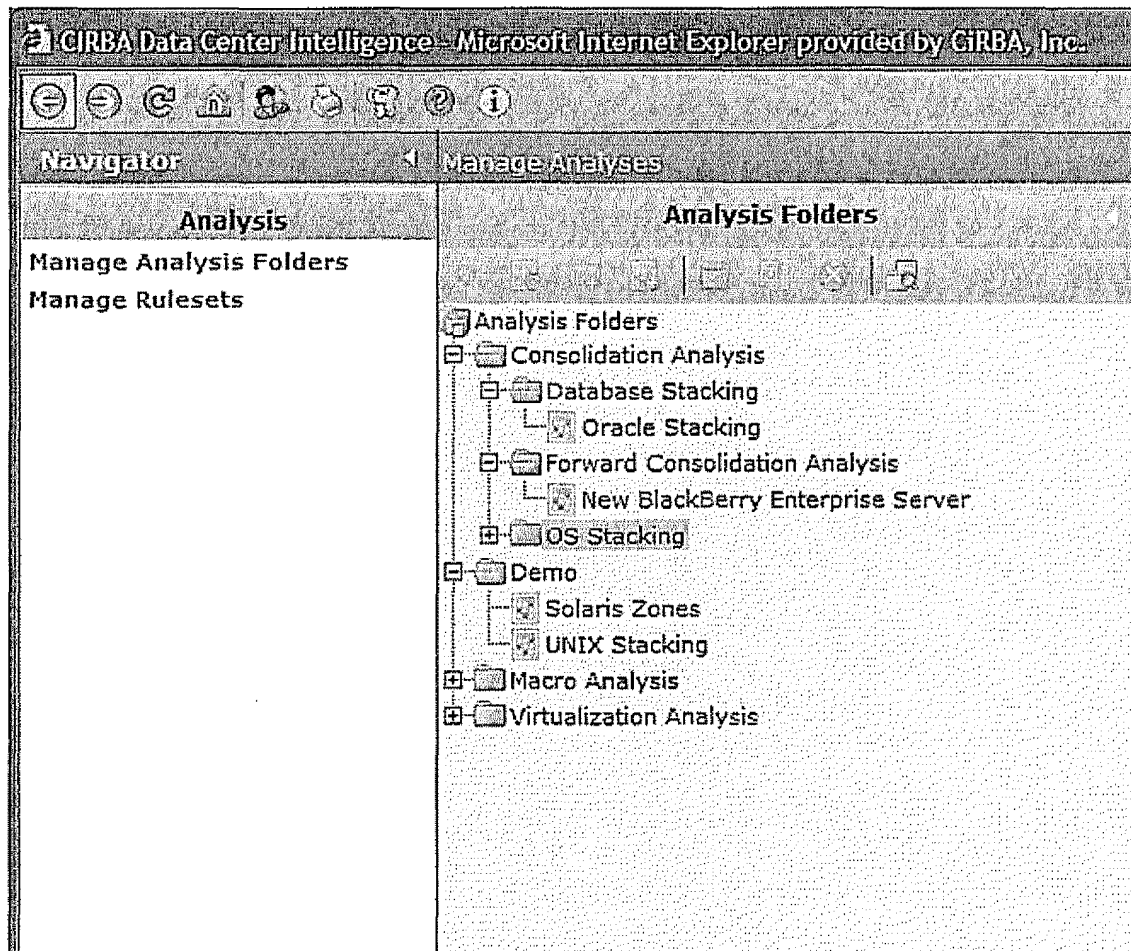


Figure 32



**Demo Stacking -- Web Page Dialog**

**Options** | **Importance Factors**

Name:

Description:

---

**System Selection** Edit

System Groups

Systems to Analyze

- dc1.unix.3809\_
- dc1.unix.3995\_
- dc1.unix.4657\_
- dc1.unix.4925\_
- dc1.unix.16588\_
- dc1.unix.16830\_

---

**Configuration Compatibility Map** Edit

Ruleset

Rules to Apply

- Configuration/Unix OS Stacking
- Business Constraints/Combined Business

---

**Workload Stacking Map** Edit

Workload Types

- CPU System
- CPU User
- Network I/O
- Memory Utilization
- Disk I/O
- Disk Space Usage

Workload Types to Apply

- CPU Utilization

Figure 33

**Ruleset Editor - Web Page Dialog**

Ruleset Category:

Ruleset Name:

Description:

Version:

Ruleset is User Defined

Ruleset Items

<input type="checkbox"/>	Item	Description	Weight	Score
<input checked="" type="checkbox"/>	OS	Different OS	100	0.00
<input checked="" type="checkbox"/>	Older OS Version	Target running older version of OS	70	0.00
<input checked="" type="checkbox"/>	Newer OS Version	Target running newer version of OS	30	0.00
<input checked="" type="checkbox"/>	Patch Level	Different Patch Levels	5	0.00
<input checked="" type="checkbox"/>	Kernel Bits	Different kernel bits	5	0.00
<input checked="" type="checkbox"/>	Time Zones	Systems are in different time zones	5	0.00
<input checked="" type="checkbox"/>	Tape Drives	Target does not have tape drive	2	0.00
<input checked="" type="checkbox"/>	Memory	Target Has Less Memory	2	0.00
<input checked="" type="checkbox"/>	Patches	Patch differences between systems	2	0.00
<input checked="" type="checkbox"/>	Installed Software	Target does not have software package installed	2	0.00
<input checked="" type="checkbox"/>	Software Versions	Different software versions	2	0.00
<input checked="" type="checkbox"/>	Abuse Kernel Parameters	Kernel configuration file does not exist or is not readable	2	0.00

Save Apply Cancel

Figure 34

**Workload Settings - Web Page Dialog**

**Audit Date Range**

☒ All available workload data

☐ Last  Days (Input value of 0 means today)

☐ Between  and  (dd/mm/yyyy)

**Only include workload data for the following days**

☒ Monday ☒ Tuesday ☒ Wednesday ☒ Thursday ☒ Friday ☒ Saturday ☒ Sunday

**Workload Type Limits**

Workload Type	Limit
CPU Utilization	<input type="text" value="90"/> %

Figure 35

**Demo Stacking - Web Page Dialog**

**Options**   **Importance Factors**

Importance Factor: 1 = Less Important, 10 = More Important

Include	Item	Item Type	Category	Importance
<input checked="" type="checkbox"/>	Unix OS Stacking	Ruleset	Configuration	5
<input checked="" type="checkbox"/>	Combined Business Constraints	Ruleset	Business Constraints	5
<input checked="" type="checkbox"/>	CPU Utilization	Workload	N/A	5

**Figure 36**



## Configuration Compatibility Report

Report Generated: Apr 21, 2007 11:47:10 AM

### Summary

Source	_dd_unix_7309_
Target	_dd_unix_11691_
Score	63
Category	Configuration
Ruleset Name	Unix OS Stacking
Version	04202007
Remediation Cost	\$ 0.00

### Summary of Differences

Description	Penalty Weight	Counts	Remediation Cost
/dev/tcp settings differ	3%	1	\$ 0.00
/dev/udp settings differ	0%	1	\$ 0.00
Kernel file group differ	1%	1	\$ 0.00
Kernel file permissions differ	1%	1	\$ 0.00
Target has different/inetD services	2%	1	\$ 0.00
Target has different/inetD services settings	2%	1	\$ 0.00
Target running nonos version of OS	30%	1	\$ 0.00
udp_max_buf settings are different	1%	1	\$ 0.00
<b>Total</b>		<b>6</b>	<b>\$ 0.00</b>

### Difference Details

Source System	Target System	Description	Module	Object	Property	Instance	Source	Target	Penalty Weight	Remediation Cost
_dd_unix_7309_	_dd_unix_11691_	/dev/tcp settings differ	CIRBA Solaris	CIRBA Solaris /DD Table (/dev/tcp)	Device Value	tcp_time_wait_interval	240000	60000	3%	\$ 0.00
_dd_unix_7309_	_dd_unix_11691_	/dev/udp settings differ	CIRBA Solaris	CIRBA Solaris /DD Table (/dev/udp)	Device Value	tcp_maxsize_multiplier	2	10		\$ 0.00
_dd_unix_7309_	_dd_unix_11691_	/dev/tcp settings differ	CIRBA Solaris	CIRBA Solaris /DD Table (/dev/tcp)	Device Value	tcp_max_buf	16384	49152		\$ 0.00
_dd_unix_7309_	_dd_unix_11691_	/dev/udp settings differ	CIRBA Solaris	CIRBA Solaris /DD Table (/dev/udp)	Device Value	tcp_recv_buf	24576	49152		\$ 0.00
_dd_unix_7309_	_dd_unix_11691_	/dev/tcp settings differ	CIRBA Solaris	CIRBA Solaris /DD Table (/dev/tcp)	Device Value	tcp_chrony_us	2	1		\$ 0.00
_dd_unix_7309_	_dd_unix_11691_	/dev/tcp settings differ	CIRBA Solaris	CIRBA Solaris /DD Table (/dev/tcp)	Device Value	tcp_wscale_always	0	1		\$ 0.00

Figure 38

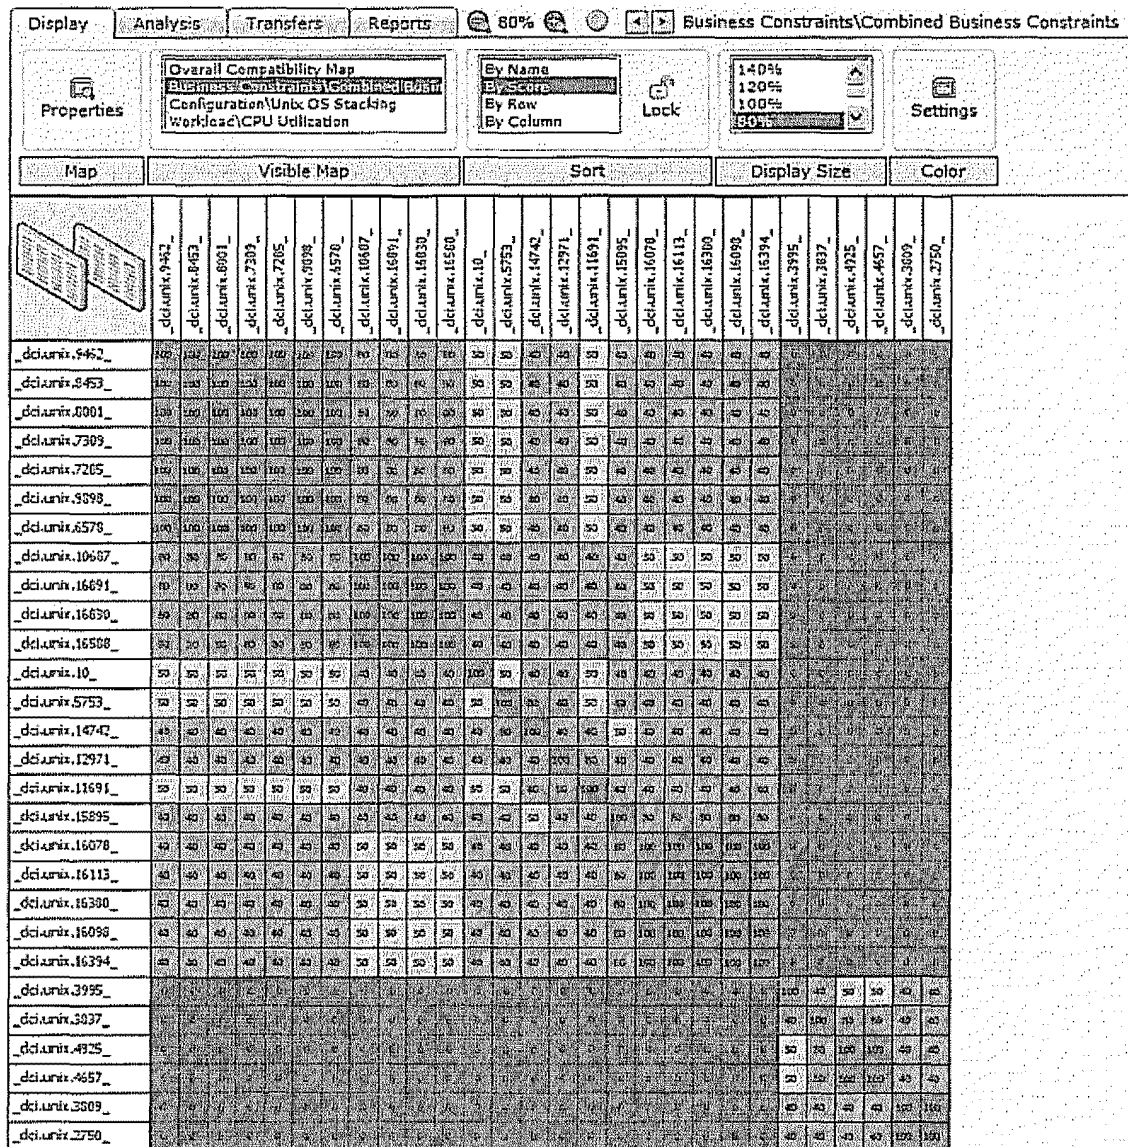


Figure 39





## Workload Stacking Report

Report Generated: Apr 21, 2007 11:44:18 AM

### Summary

Source	_dcl.unix.16113_
Target	_dcl.unix.16076_
Score	78
Workload Type	CPU Utilization
Usage Limit	90 %
Benchmark Type	Cpu
Source Benchmark	2972
Target Benchmark	1396

### Scorecard

Criteria	Value	Score
Peak 4th Quartile (Like Times)	155	28
Peak 4th Quartile (Worst Case)	161	21
Peak 3rd Quartile (Like Times)	9	100
Peak 3rd Quartile (Worst Case)	10	100
Aggregate		78

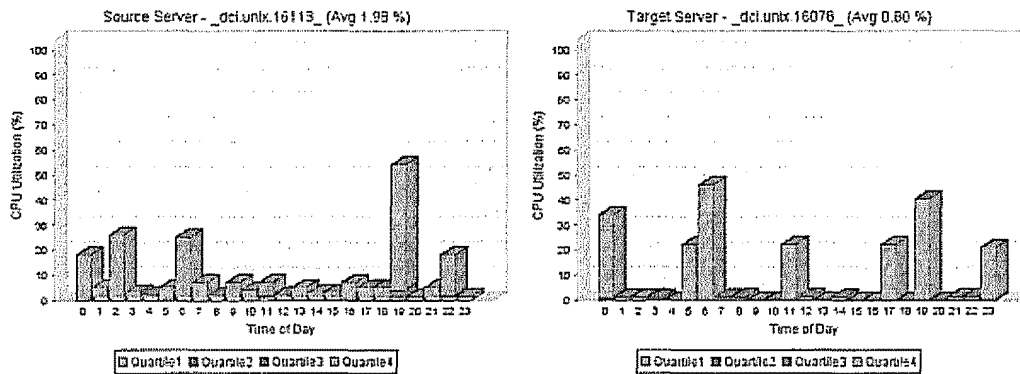




Figure 41





CPU Utilization peaks at 161 % between the hours of 6 AM and 7 AM when the worst times are stacked together.



#### Audit Information

Type	System	Audit	Audit Timestamp	Family	P
Source	_dc1.unix.16113_	Solaris9-2	Jan 11, 2007 12:00:00 AM	Solaris	_dc1.unix.11049
Target	_dc1.unix.16078_	Solaris9-2	Jan 11, 2007 12:00:00 AM	Solaris	_dc1.unix.11040

Figure 42


Display	Analysis	Transfers			Reports			100%			Overall Compatibility Map																	
	_dcl.unix.7309_	_dcl.unix.8001_	_dcl.unix.8453_	_dcl.unix.9462_	_dcl.unix.10687_	_dcl.unix.7285_	_dcl.unix.6578_	_dcl.unix.9898_	_dcl.unix.5753_	_dcl.unix.10_	_dcl.unix.14742_	_dcl.unix.16394_	_dcl.unix.16113_	_dcl.unix.16078_	_dcl.unix.16098_	_dcl.unix.16380_	_dcl.unix.15895_	_dcl.unix.16830_	_dcl.unix.16891_	_dcl.unix.16588_	_dcl.unix.11691_	_dcl.unix.12971_	_dcl.unix.3995_	_dcl.unix.3837_	_dcl.unix.4657_	_dcl.unix.4925_	_dcl.unix.3809_	_dcl.unix.2750_
_dcl.unix.7309_	100	85	81	61	69	64	65	59	35	32	25	24	26	25	25	21	21	47	36	45	22	19	0	0	0	0	0	0
_dcl.unix.8001_	94	100	88	61	69	67	67	61	36	33	26	25	26	26	26	26	25	51	41	16	27	21	0	0	0	0	0	0
_dcl.unix.8453_	89	89	100	61	69	71	71	65	37	36	26	25	26	26	26	26	26	52	45	47	29	27	0	0	0	0	0	0
_dcl.unix.9462_	91	91	91	100	72	79	78	74	39	37	26	26	26	27	27	26	27	53	50	49	32	25	0	0	0	0	0	0
_dcl.unix.10687_	74	74	74	50	100	62	60	55	30	29	26	32	33	33	33	33	27	66	61	61	26	25	0	0	0	0	0	0
_dcl.unix.7285_	95	65	20	0	36	100	5	2	15	1	0	17	26	39	19	13	7	28	6	30	0	5	0	0	0	0	0	0
_dcl.unix.6578_	85	28	0	0	17	0	100	0	7	0	19	13	1	0	21	1	0	12	0	21	0	4	0	0	0	0	0	0
_dcl.unix.9898_	83	60	19	0	42	0	0	100	1	1	17	11	24	16	10	14	7	29	0	21	0	0	0	0	0	0	0	0
_dcl.unix.5753_	38	0	0	0	0	0	0	0	100	0	29	7	27	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
_dcl.unix.10_	42	13	0	0	0	0	0	0	0	100	29	16	13	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
_dcl.unix.14742_	13	2	0	0	0	0	0	0	0	0	100	27	36	2	7	3	1	5	0	0	0	0	0	0	0	0	0	0
_dcl.unix.16394_	10	4	0	0	0	0	0	0	0	0	26	100	82	30	30	15	13	11	0	12	8	0	0	0	0	0	0	0
_dcl.unix.16113_	11	0	8	0	14	0	0	0	0	0	36	97	100	71	69	64	54	32	31	36	10	17	0	0	0	0	0	0
_dcl.unix.16078_	11	11	14	0	13	0	0	0	0	0	38	92	95	100	99	88	75	41	33	43	0	43	0	0	0	0	0	0
_dcl.unix.16098_	14	11	14	0	14	0	0	0	0	0	38	88	95	99	100	88	74	41	33	43	0	43	0	0	0	0	0	0
_dcl.unix.16380_	17	11	14	0	14	0	0	0	0	0	38	90	97	92	91	100	70	40	31	39	10	18	0	0	0	0	0	0
_dcl.unix.15895_	11	13	10	0	14	0	0	0	0	0	46	68	73	72	71	67	100	31	23	29	0	0	0	0	0	0	0	0
_dcl.unix.16830_	23	21	19	0	27	0	14	10	0	0	37	41	46	41	41	39	31	100	61	77	0	11	0	0	0	0	0	0
_dcl.unix.16891_	27	27	20	0	28	12	13	11	0	0	36	41	46	41	41	39	32	77	100	73	0	11	0	0	0	0	0	0
_dcl.unix.16588_	40	0	0	0	0	0	0	0	0	0	34	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0
_dcl.unix.11691_	34	1	13	13	13	13	13	13	12	12	11	12	11	11	11	11	11	10	10	100	75	0	0	0	0	0	0	0
_dcl.unix.12971_	11	10	10	13	13	0	0	0	20	5	14	16	11	11	11	11	10	11	10	0	74	100	0	0	0	0	0	0
_dcl.unix.3995_	16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	100	74	17	44	21	22	
_dcl.unix.3837_	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	35	100	72	69	35	23	
_dcl.unix.4657_	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	46	48	100	82	0	23	0
_dcl.unix.4925_	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	89	100	0	0	0	0
_dcl.unix.3809_	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	35	34	36	35	100	63	0
_dcl.unix.2750_	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	35	40	36	33	62	100	0

Figure 43

# Overall Compatibility Report

Report Generated: Apr 21, 2007 11:48:34 AM

## Summary

Source	dc1.unix.8453
Target	dc1.unix.9462
Overall Score	61

Category	Name	Importance	Score	Remediation Cost
Business Constraints	Combined Business Constraints	5	100	\$ 0.00
Configuration	Unix OS Stacking	5	88	\$ 0.00
Workload	CPU Utilization	5	70	
Total				\$ 0.00

## System Configuration Compatibility

### Combined Business Constraints Ruleset Summary ^

Overall Score: 100

Description	Penalty Weight	Count	Remediation Cost
Total	100	0	\$ 0.00

### Unix OS Stacking Ruleset Summary ^

Overall Score: 88

Description	Penalty Weight	Count	Remediation Cost
Different software versions	2%	1	\$ 0.00
Kernel configuration file does not exist on target	2%	1	\$ 0.00
Kernel file group differ	1%	1	\$ 0.00
Patch differences between systems	2%	1	\$ 0.00
Target Has Less Memory	2%	1	\$ 0.00
Target does not have software package installed	2%	1	\$ 0.00
Target has different InetD services	2%	1	\$ 0.00
Target has less allowable processes	0%	1	\$ 0.00
Total	88	8	\$ 0.00

## Figure 44

# Workload Stacking Compatibility

## Workloads Summary

Workload Type	Usage Limit	Benchmark	Source Benchmark	Target Benchmark
CPU Utilization	90 %	cpi	940	161

## Workloads Scorecard

	Peak 4th Quartile (Like Times)		Peak 4th Quartile (Worst Case)		Peak 3rd Quartile (Like Times)		Peak 3rd Quartile (Worst Case)		Aggregate Workload
	Value	Score	Value	Score	Value	Score	Value	Score	
CPU Utilization	374	0	374	0	29	100	30	100	78
Overall									70

## CPU Utilization Workload \*




Figure 45

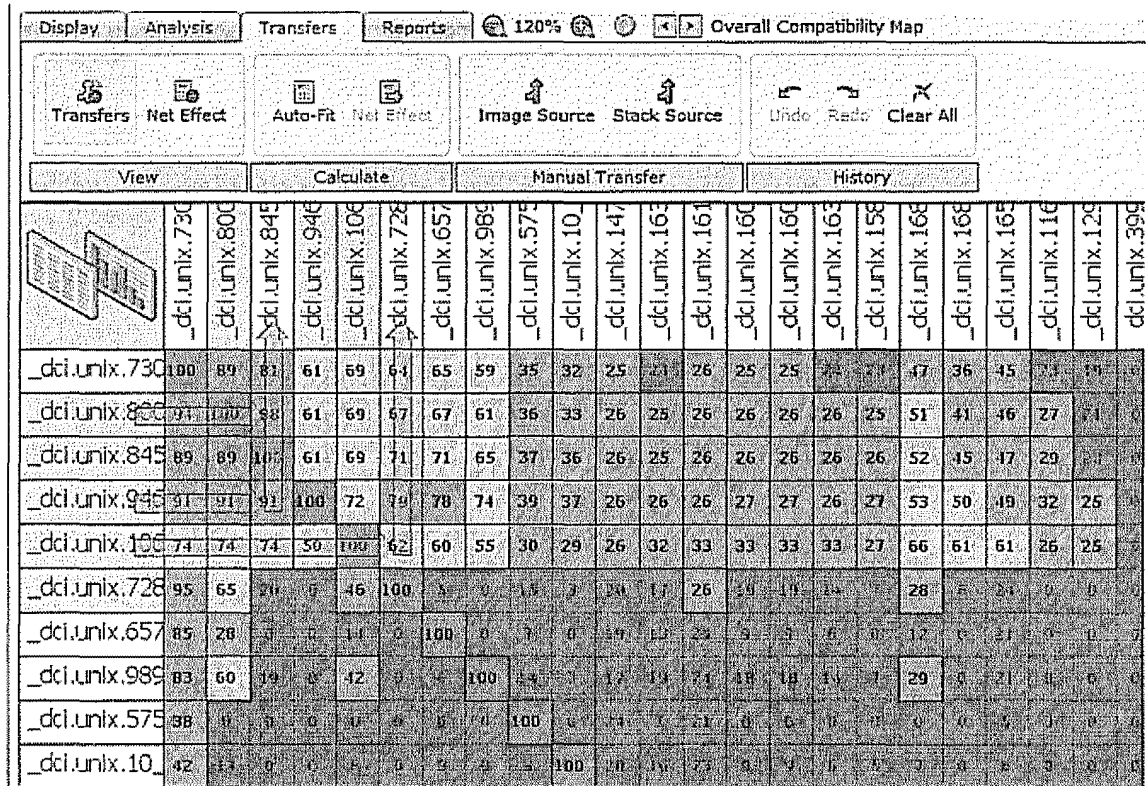


Figure 46



## Overall Compatibility Report

Report Generated: Apr 21, 2007 11:54:26 AM

### Summary

Transfer Source	dcl.unix.8462
Transfer Source	dcl.unix.8001
Target	dcl.unix.8453
Overall Score	91

Category	Name	Importance	Score	Remediation Cost
Business Constraints	Combined Business Constraints	5	100	\$ 0.00
Configuration	Unix OS Stacking	5	89	\$ 0.00
Workload	CPU Utilization	5	95	\$ 0.00
Total				\$ 0.00

### System Configuration Compatibility

#### Combined Business Constraints Ruleset Summary \*

Overall Score: 100

Description	Penalty Weight	Count	Remediation Cost
Total	100	0	\$ 0.00

#### Unix OS Stacking Ruleset Summary \*

Overall Score: 89

Description	Penalty Weight	Count	Remediation Cost
Different software versions	2%	1	\$ 0.00
Kernel file group differ	1%	1	\$ 0.00
Patch differences between systems	2%	1	\$ 0.00
Target Has Less Memory	2%	1	\$ 0.00
Target does not have software package installed	2%	1	\$ 0.00
Target has different InetD services	2%	1	\$ 0.00
Total	89	6	\$ 0.00

Figure 48

**Auto-Fit Transfers -- Web Page Dialog**

**Transfer Type**

☒ Stacking  
☐ Virtualization

**Optimization**

☐ Quick Fit  
☒ Detailed Analysis

**Limits**

☒ Minimum Source Systems per Target: 1

☒ Maximum Systems per Target: 5

Do not transfer if Target score is below: 75

☒ Clear existing Transfers

OK Cancel

**Figure 49**



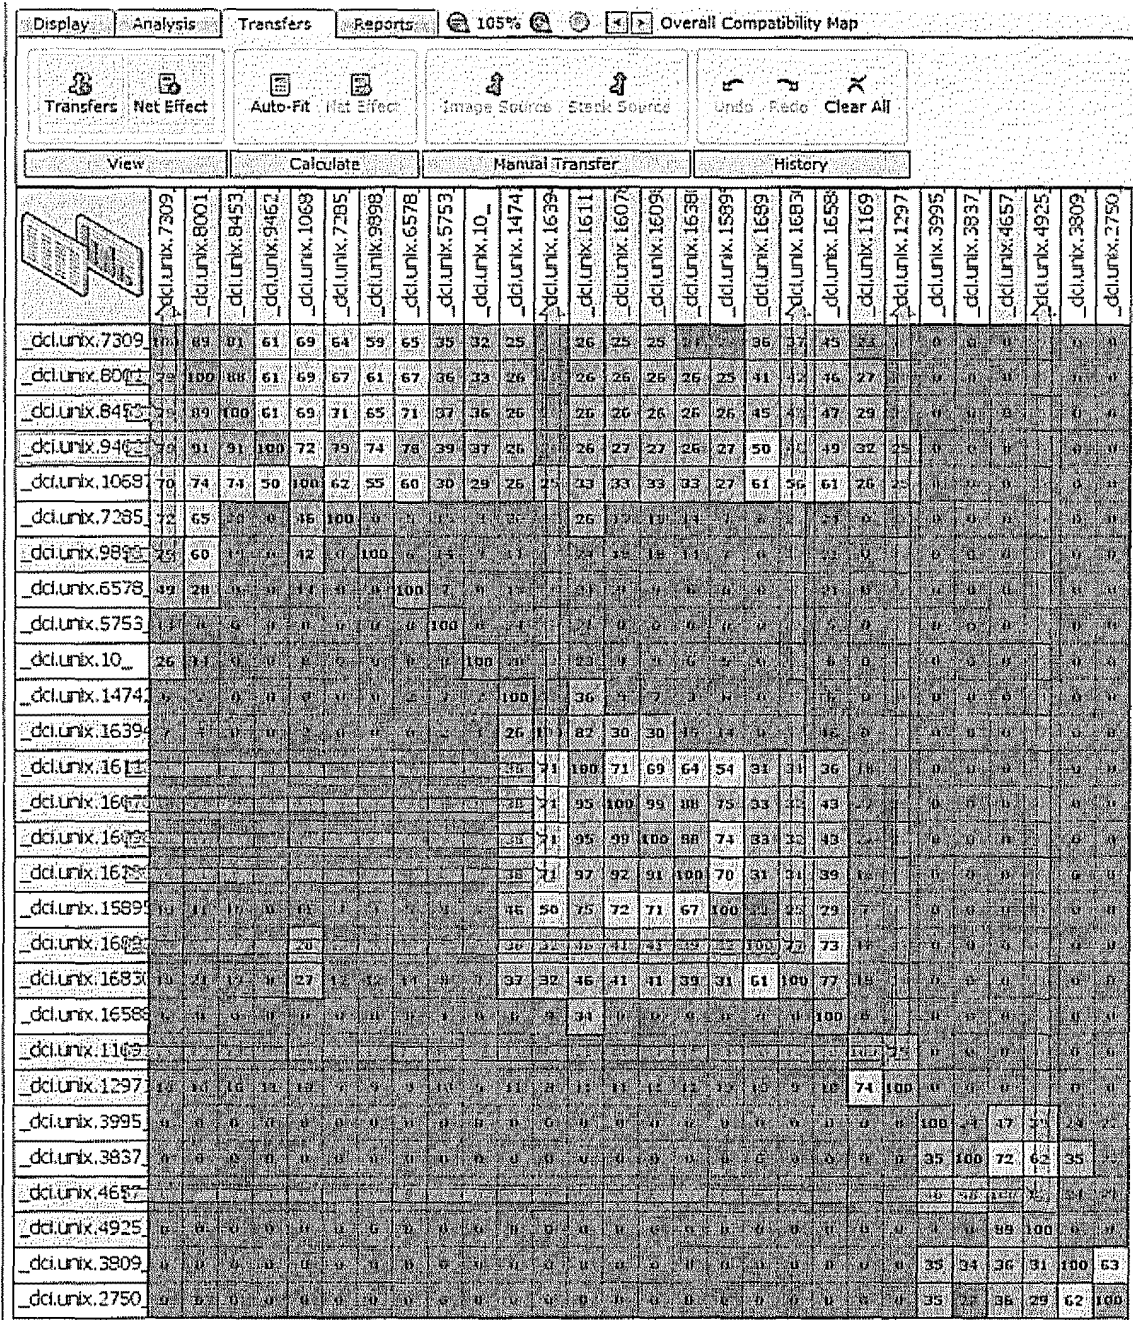


Figure 50

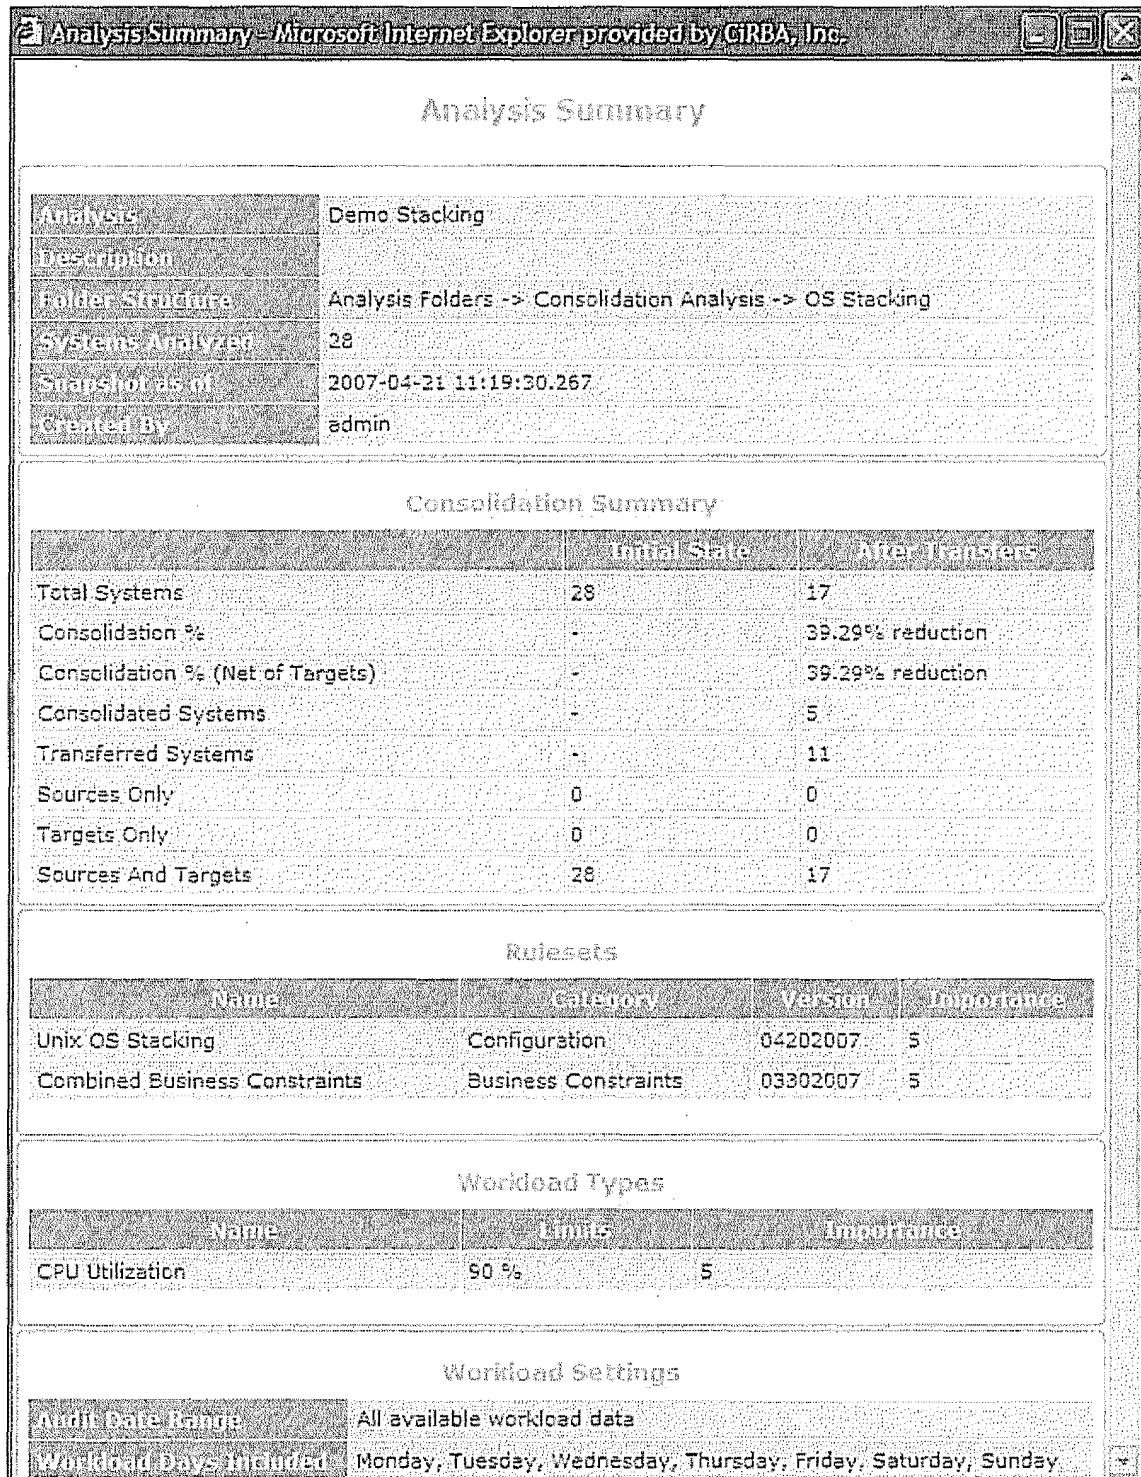


Figure 51

Transfers		
Source	Transfer	Target
_dci.unix.9898_	Stacks onto	_dci.unix.7309_
_dci.unix.8453_	Stacks onto	_dci.unix.7309_
_dci.unix.9462_	Stacks onto	_dci.unix.7309_
_dci.unix.8001_	Stacks onto	_dci.unix.7309_
_dci.unix.16113_	Stacks onto	_dci.unix.16394_
_dci.unix.16098_	Stacks onto	_dci.unix.16394_
_dci.unix.16380_	Stacks onto	_dci.unix.16394_
_dci.unix.16078_	Stacks onto	_dci.unix.16394_
_dci.unix.4657_	Stacks onto	_dci.unix.4925_
_dci.unix.16891_	Stacks onto	_dci.unix.16830_
_dci.unix.11691_	Stacks onto	_dci.unix.12971_

**Figure 52**