

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

2
3 **CROSS-REFERENCE TO RELATED APPLICATIONS**
4

5 **[0001]** This application is a continuation of U.S. Patent Application No. 14/341,471 filed on
6 July 25, 2014, which is a continuation of U.S. Patent Application No. 11/738,936 filed on April
7 23, 2007, which is a continuation-in-part of U.S. Patent Application No. 11/535,355 and also a
8 continuation-in-part of U.S. Patent Application No. 11/535,308 both filed on September 26,
9 2006, both of which claim priority from U.S. Provisional Patent Application No. 60/745,322 filed
10 April 21, 2006, all incorporated herein by reference.

11 **TECHNICAL FIELD**
12

13 **[0002]** The present invention relates to information technology infrastructures and has
14 particular utility in determining compatibility of computer systems in such infrastructures.

15 **BACKGROUND**

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

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

29 **[0005]** While the benefits of a distributed approach are numerous and well understood,
30 there has arisen significant practical challenges in managing such systems for optimizing
31 efficiency and to avoid redundancies and/or under-utilized hardware. In particular, one

1 challenge occurs due to the sprawl that can occur over time as applications and servers
2 proliferate. Decentralized control and decision making around capacity, the provisioning of new
3 applications and hardware, and the perception that the cost of adding server hardware is
4 generally inexpensive, have created environments with far more processing capacity than is
5 required by the organization.

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

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

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

25 **[0009]** OS-Level application stacking involves moving the applications and data from one or
26 more systems to the consolidated system. This can effectively replace multiple operating
27 system instances with a single OS instance, e.g. system A running application X and system B
28 running application Y are moved onto system C running application Z such that system C runs

1 applications X, Y and Z, and system A and B are no longer required. This strategy is applicable
2 to all operation system types, e.g. Windows™, Linux™, Solaris™, AIX™, HP-UX™, etc.

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

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

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

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

23 **[0014]** It is therefore an object of the following to obviate or mitigate the above-described
24 disadvantages.

25 SUMMARY

26 **[0015]** In one aspect, a method for determining a consolidation solution for a plurality of
27 computer systems is provided comprising obtaining a data set comprising a compatibility score

1 for each pair of the plurality of computer systems, each the compatibility score being indicative
2 of the compatibility of one of the plurality of computer systems with respect to another of the
3 plurality of computer systems; determining one or more candidate transfer sets each indicating
4 one or more of the computer systems capable of being transferred to a target computer system;
5 selecting a desired one of the one or more candidate transfer sets; and providing the desired
6 one as a consolidation solution.

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

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

22 BRIEF DESCRIPTION OF THE DRAWINGS

23 **[0018]** An embodiment of the invention will now be described by way of example only with
24 reference to the appended drawings wherein:

25 **[0019]** Figure 1 is a block diagram of an analysis program for evaluating the compatibility of
26 computer systems to identify consolidation solutions.

27 **[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.
- 21 **[0038]** Figure 20 is a flow diagram of workload data extraction process.

- 1 **[0039]** Figure 21 is a flow diagram of the 1-to-1 workload compatibility analysis.
- 2 **[0040]** Figure 22 is an example of the workload compatibility analysis result details.
- 3 **[0041]** Figure 23 is an example of the overall compatibility analysis result details.
- 4 **[0042]** Figure 24(a) is a high level process flow diagram of the multi-dimensional
5 compatibility analysis.
- 6 **[0043]** Figure 24(b) is a flow diagram showing the multi-dimensional analysis.
- 7 **[0044]** Figure 24(c) is a flow diagram showing use of a rule set in an N-to-1 compatibility
8 analysis.
- 9 **[0045]** Figure 24(d) is a flow diagram showing use of a rule set in an N-by-N compatibility
10 analysis.
- 11 **[0046]** Figure 25 is a process flow diagram of the multi-dimensional workload compatibility
12 analysis.
- 13 **[0047]** Figure 26 is an example multi-dimensional compatibility analysis map.
- 14 **[0048]** Figure 27 is an example of the multi-dimensional compatibility analysis result details
15 for a rule set.
- 16 **[0049]** Figure 28 is an example of the multi-dimensional workload compatibility analysis
17 result details.
- 18 **[0050]** Figure 29 is a process flow diagram of the consolidation analysis.
- 19 **[0051]** Figure 30 is a process flow diagram of an auto fit algorithm used by the consolidation
20 analysis.
- 21 **[0052]** Figure 31 is an example of a consolidation solution produced by the consolidation
22 analysis.
- 23 **[0053]** Figure 32 shows an example hierarchy of analysis folders and analyses.

- 1 **[0054]** Figure 33 shows the screen for creating and editing the analysis input parameters.
- 2 **[0055]** Figure 34 shows an example of the rule set editor screen.
- 3 **[0056]** Figure 35 shows an example of the screen for editing workload settings.
- 4 **[0057]** Figure 36 shows an example screen to edit the importance factors used to compute
5 the overall compatibility score.
- 6 **[0058]** Figure 37 is an example 1-to-1 compatibility map.
- 7 **[0059]** Figure 38 is example configuration compatibility analysis details.
- 8 **[0060]** Figure 39 is an example 1-to-1 compatibility map for business constraints.
- 9 **[0061]** Figure 40 is an example 1-to-1 workload compatibility map.
- 10 **[0062]** Figure 41 is an example workload compatibility report.
- 11 **[0063]** Figure 42 is an example workload details report with stacked workloads.
- 12 **[0064]** Figure 43 is an example of a 1-to-1 overall compatibility map.
- 13 **[0065]** Figure 44 is an example of the 1-to-1 overall compatibility details report.
- 14 **[0066]** Figure 45 shows the workload details of the overall compatibility report.
- 15 **[0067]** Figure 46 shows example transfers on a compatibility map with net effect off.
- 16 **[0068]** Figure 47 shows example transfers on a compatibility map with net effect on.
- 17 **[0069]** Figure 48 is an example multi-dimensional compatibility details report.
- 18 **[0070]** Figure 49 shows an example of the consolidation analysis (auto fit) input screen.
- 19 **[0071]** Figure 50 is an example overall compatibility map for the consolidation solution.
- 20 **[0072]** Figure 51 is an example consolidation summary report.

1 **[0073]** Figure 52 shows the example transfers that comprise the consolidation solution.

2 DETAILED DESCRIPTION OF THE DRAWINGS

3 **Analysis Program Overview**

4 **[0074]** A block diagram of an analysis program 10 for determining compatibilities in
5 computing environment 12 is provided in Figure 1. The analysis program 10, accessed through
6 a computer station 14, gathers data 18 pertaining to a collection of systems to be consolidated
7 16. The analysis program 10 uses the gathered data 18 to evaluate the compatibility of the
8 computer systems 28 and provide a roadmap 20 specifying how the original set of systems can
9 be consolidated to a smaller number of systems 22.

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

12 **System Data Parameters**

13 **[0076]** A distinct data set is obtained for each system 16 to contribute to the combined
14 system data 18 shown in Figure 2. Each data set comprises one or more parameters that relate
15 preferably to technical 24, business 26 and workload 28 characteristics or features of the
16 respective system 16. The parameters can be evaluated by scrutinizing program definitions,
17 properties, objects, instances and any other representation or manifestation of a component,
18 feature or characteristic of the system 16. In general, a parameter is anything related to the
19 system 16 that can be evaluated, quantified, measured, compared etc.

20 **[0077]** Examples of technical parameters relevant of the consolidation analysis include the
21 operating system, OS version, patches, application settings, hardware devices, etc.

22 **[0078]** Examples of business parameters of systems relevant to the consolidation analysis
23 include the physical location, organization department, data segregation requirements, owner,
24 service level agreements, maintenance windows, hardware lease agreements, software
25 licensing agreements, etc.

1 [0079] Examples of workload parameters relevant to consolidation analysis include various
2 resource utilization and capacity metrics related to the system processor, memory, disk storage,
3 disk I/O throughput and network bandwidth utilization.

4 **System and Entity Models**

5 [0080] The system data parameters associated with a system 16 comprise the system
6 model used in the analyses.

7 [0081] In the following examples, a source system refers to a system from which
8 applications and/or data are to be moved, and a target server or system is a system to which
9 such applications and/or data are to be moved. For example, an underutilized environment
10 having two systems 16 can be consolidated to a target system (one of the systems) by moving
11 applications and/or data from the source system (the other of the systems) to the target system.

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

20 [0083] Hypothetical systems can be created through data imports, cloning from actual
21 systems models, and manual specification by users, etc. The system model can be minimal
22 (sparse) or include as much data as an actual system model. These system models may also
23 be further modified to address the analysis requirements.

24 [0084] The compatibility analysis can also be generalized to evaluate entities beyond
25 physical, virtual or hypothetical systems. For example, entities can be components that
26 comprise systems such as applications and database instances. By analysing the compatibility
27 of database instances and database servers with database stacking rule sets, database
28 consolidation can also be assessed. Similarly, application consolidation can be evaluated by

analyzing application servers and instances with application stacking rules. The entity could also be a logical application system and technical data can pertain to functional aspects and specifications of the entity.

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

Consolidation and Transfers

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

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

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

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

1 **[0091]** The following discusses compatibilities between systems 16 based on the
2 parameters to determine if efficiencies can be realized by consolidating either entire systems 16
3 or aspects or components thereof.

4 **[0092]** The analyses employ differential rule sets 28 to evaluate and quantify the
5 compatibility of systems 16 with respect to technical configuration and business related factors
6 comprised in the gathered system data 18. Similarly, workload compatibility of a set of systems
7 16 is assessed using workload stacking and scoring algorithms 30. The results of configuration,
8 business and workload compatibility analyses are combined to produce an overall compatibility
9 score for a set of systems 16.

10 **[0093]** In addition to compatibility scores, the analysis provides details that account for the
11 actual scores. The scores can be presented in color coded maps 32, 34 and 36 that illustrate
12 patterns of the compatibility amongst the analyzed systems as shown in Figures 4, 5 and 6
13 respectively.

14 **Analysis Modes**

15 **[0094]** A collection of systems 16 to be consolidated can be analyzed in one of three
16 modes: 1-to-1 compatibility, multi-dimensional compatibility and consolidation analyses. These
17 analyses share many common aspects but can be performed independently.

18 **[0095]** The 1-to-1 compatibility analysis evaluates the compatibility of every possible
19 source-target pair combination in the collection of systems 16 on a 1-to-1 basis. This analysis is
20 useful in assessing single transfer consolidation candidates. In practice, it may be prudent to
21 consolidate systems 16 incrementally and assess the impact of each transfer before proceeding
22 with additional transfers.

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

1 **[0097]** The consolidation analysis searches for a consolidation solution that minimizes the
2 number of remaining source and target entities after the proposed transfers are applied, while
3 meeting requisite compatibility constraints. This analysis employs the multi-dimensional
4 compatibility analysis described above to evaluate the compatibility of postulated transfer sets.

5 **[0098]** The analysis program 10 performs consolidation analyses for virtualization and
6 stacking strategies as will be explained in greater detail below, however, it will be appreciated
7 that other consolidation strategies may be performed according to similar principles.

8 **Analysis Program Architecture**

9 **[0099]** A block diagram of the analysis program 10 is shown in Figure 7. The flow of data
10 18 through the program 10 begins as an audit engine 40 pulls audit data 18 from audited
11 environments 42. The data works its way up to the web client 44 which displays an output on a
12 user interface, e.g. on computer 14. The program 10 is preferably a client-server application
13 that is accessed via the web client 44.

14 **[00100]** An audit engine 40 communicates over one or more communication protocols with
15 audited environments 42 comprised of the actual systems 16 being analysed. The audit engine
16 40 typically uses data acquisition adapters to communicate directly with the end points (e.g.
17 servers) or through software systems that manage the end points (e.g. management
18 frameworks 46 and/or agent instrumentation 48 and/or direct access 50).

19 **[00101]** Alternatively, system data 18 can be imported from third party tools, e.g. inventory
20 applications, performance monitoring tools etc., or can be obtained through user data entry.
21 Examples of such third-party data having file formats that can be imported include comma
22 separated values (CSV), extensible markup language (XML) and well formatted text files such
23 as those generated by the UNIX™ system activity reporter (SAR).

24 **[00102]** The audit engine 40, uses a set of audit request templates 52 that define the data 18
25 to acquire from the systems 16. Once collected, the data 18 is stored in an audit data repository
26 54. Data 18 referenced by the analysis rule sets 28 and required by the workload analysis 30
27 are extracted and stored in separate data caches 56 and 58 respectively.

1 **[00103]** Aliases 60, differential rule sets 28, workload data types 30 and benchmark
2 specifications comprise some of the analysis-related metadata 62 definitions used by the
3 program 10. Aliases 60 extract and normalize system data 18 from a variety of data sources to
4 a common model for analysis. Rule sets 28 examine system compatibility with respect to
5 technical configuration and business-related factors. The workload definitions 30 specify the
6 system resource parameters and benchmarks for analyzing workload compatibility.

7 **[00104]** An analysis engine 64 is also provided, which comprises compatibility and
8 consolidation analysis engines 66 and 68 respectively. The compatibility analysis evaluates the
9 compatibility of systems 16 through rule sets 28 and workload stacking algorithms 30. The
10 consolidation analysis engine 68 leverages the compatibility analysis and employs constraint-
11 based optimization algorithms to find consolidation solutions that allows the environment 12 to
12 operate with fewer systems 16.

13 **[00105]** The program 10 has a report engine 70 that utilizes report templates 72 for
14 generating reports that convey the analysis results. Typically, the program 10 includes a web
15 interface layer 74 that allows web client 44 users to enter settings, initiate an audit or analysis,
16 view reports etc.

17 **Analysis Data Sources**

18 **[00106]** The audit data 18 can be acquired using tools such as the table 76 shown in Figure 8
19 that illustrate the various types of configuration settings that are of interest and from which
20 sources they can be obtained. Figure 8 also provides a mapping to where the sample workload
21 data can be obtained. In Figure 8, a number of strategies 78 and sub-strategies 80 map to
22 various configuration and workload sources, collectively referred to by numeral 82. As
23 discussed making reference to Figure 8, the strategies 78 may relate to database consolidation,
24 OS-level stacking, application server stacking, virtualization, and many others. Each strategy
25 78 includes a set of sub-strategies 80, which in turn map to specific rule sets 28. The rule sets
26 28, which will be explained in greater detail below, determine whether or not a particular setting
27 or system criterion/criteria have been met and thus how different one system 16 is to the next.
28 The rule sets 28 can also indicate the cost of remediating such differences.

1 **[00107]** The table 76 lists the supported consolidation strategies and the relevant data
2 sources that should be audited to perform the corresponding consolidation analysis. In general,
3 collecting more basis data 18 improves the analysis results. The table 76 enables the analysis
4 program 10 to locate the settings and information of interest based on the strategy 78 or sub-
5 strategy 80 (and in turn the rule set 28) that is to be used to evaluate the systems 16 in the
6 environment 12. The results can be used to determine source/target candidates for analysing
7 the environment for the purpose of, e.g. consolidation, compliance measures etc.

8 **Analysis Process Overview**

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

14 **[00109]** In process A, the system data 18 collected via audits or imports as discussed above
15 is prepared for use by the analyses. The compatibility and consolidation analyses processes B,
16 C and D can be performed independently. The analyses share a common analysis input
17 specification and get system data 18 from the data repository 54 and caches 56 and 58. The
18 multi-dimensional compatibility and consolidation analyses take additional inputs in the form of a
19 consolidation solution and auto fit input parameters 84 and 86 respectively.

20 **[00110]** The 1-to-1 compatibility analysis process B evaluates the compatibility of each
21 system pair on a 1-to-1 basis. In contrast, the multi-dimensional analysis process C evaluates
22 the compatibility of each transfer set 23 in the consolidation solution that was specified as part
23 of its input.

24 **[00111]** The consolidation analysis process D searches for the best consolidation solution
25 that fulfills the constraints defined by the auto fit input 86. The consolidation analysis employs
26 the multi-dimensional compatibility analysis C to assess potential transfer set candidates.

27 **Data Load and Extraction**

1 [00112] A process flow diagram for the data load and extraction process A is illustrated in
2 Figure 10. System data including technical configuration, business related and workload
3 collected through audits, data import and user input are prepared for use by the analyses
4 processes B, C and D.

5 [00113] When system data 18 and attributes are loaded into the analysis program 10, they
6 are stored in the audit data repository 54 and system attribute table 55, respectively. As well,
7 system data 18 referenced by rule set items 28, workload types 30 and benchmarks are
8 extracted and loaded into their respective caches 56, 58. Alias specifications 60 describe how
9 data can be extracted and if necessary, normalized from a variety of data sources.

10 [00114] The data repository 54 and caches 56 and 58 thus store audited data 18, system
11 attributes, the latest rule set data, historical workload data and system workload benchmarks.

12 **1-to-1 Compatibility Analysis**

13 [00115] A high level flow diagram of the 1-to-1 compatibility analysis is shown in Figure 11.
14 The 1-to-1 compatibility analysis can take into account analysis input, including input regarding
15 the systems 16 to be analyzed, rule set related parameters, workload related parameters,
16 workload benchmarks and importance factors 88 used to compute overall scores.

17 [00116] The compatibility analysis evaluates the compatibility of every specified system as
18 source-target pairs on a 1-to-1 basis. This analysis produces a compatibility score for each
19 system pair so that analyzing a collection of ten (10) systems 16 produces 10x10 scores. The
20 compatibility analysis is based on the specified rule sets and workload types.

21 [00117] An analysis may be based upon zero or more rule sets and zero or more workload
22 types, such that at least one rule set or workload type is selected. Example rule sets 28 and
23 corresponding descriptions are shown in Figure 12, and example workload types 30 and
24 corresponding descriptions are shown in Figure 13.

25 [00118] The selection of rule sets 28 and workload types 30 for an analysis depends on the
26 systems 28 and the consolidation strategy to analyze. For example, to assess the consolidation
27 of a set of UNIX™ systems 16, an analysis may employ the UNIX™ application stacking,

location, maintenance window and ownership rule sets 28, and CPU, memory, disk space, disk I/O and network I/O workload types 30.

1-to-1 Compatibility Analysis Process Flow

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

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

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

[00122] Upon completion of the compatibility analysis, the results 20 are provided to the user. 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

1 Application No. 11/535,308 filed on September 26, 2006, and entitled "Method for Evaluating
2 Computer Systems", the contents of which are incorporated herein by reference.

3 **[00125]** With respect to the following description of the rule sets 28 and the general
4 application of the rule sets 28 for detecting system incompatibilities by evaluating differences
5 between data parameters of systems 16, the following alternative nomenclature may be used.
6 A target system refers to a system being evaluated, and a baseline system is a system to which
7 the target system is being compared. The baseline and target systems may be the same
8 system 16 at different instances in time (baseline = prior, target = now) or may be different
9 systems 16 being compared to each other. As such, a single system 16 can be evaluated
10 against itself to indicate changes with respect to a datum as well as how it compares to its
11 peers. It will be appreciated that the terms "source system" and "baseline system" are herein
12 generally synonymous, whereby a source system is a type of baseline system.

13 **[00126]** Figure 1 illustrates the relationships between system data 18 and the analysis
14 program 10. Data 18 is obtained from the source and target computer systems 16 and is used
15 to analyze the compatibility between the systems 16. In this example, the parameters are
16 evaluated to determine system compatibilities for a consolidation strategy. A distinct data set 18
17 is preferably obtained for each system 16 (or instance in time for the same system 16 as
18 required).

19 **[00127]** Rule sets 28 are computer readable and storable so that they may be accessed by
20 the program 10 and modified if necessary, for use in evaluating the computer systems 16.

21 **[00128]** Rule sets 28 are groupings of rules that represent higher-level considerations such
22 as business objectives or administrative concerns that are taken into account when reporting on
23 or analysing the systems 16. In Figure 15, six rules 43, A, B C, D, E and F are grouped into
24 three rule sets 28, Rule Set 1, 2 and 3. It will be appreciated that there may be any number of
25 rules in any number of rule sets 28 and those shown in Figure 15 are for illustrative purposes
26 only.

27 **[00129]** Rules evaluate data parameters according to rule definitions to determine
28 incompatibilities due to differences (or contentious similarities) between the baseline and target
29 systems. The rule definitions include penalty weights that indicate the importance of the

1 incompatibility as they relate to the operation of the systems 16. The penalty weights are
2 applied during an evaluation if the incompatibility is detected. The evaluation may include the
3 computation of a score or generation of other information indicative of nature of the
4 incompatibilities between the baseline and target systems.

5 **[00130]** Rules comprised by a rule set 28 may reference common parameters but perform
6 different tests to identify different forms of incompatibilities that may have different levels of
7 importance. For example a version four operating system versus a version three operating
8 system may be considered less costly to remedy and thus less detrimental than a version five
9 operating system compared to a version one operating system. As can be seen, even though
10 the operating systems are different in both cases, the nature of the difference can also be
11 considered and different weights and/or remedies applied accordingly.

12 **[00131]** Rules can also test for similarities that indicate contentions which can result in
13 incompatibilities between systems. For example, rules can check for name conflicts with
14 respect to system names, database instance names, user names, etc.

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

22 **Rule Set Specification**

23 **[00133]** Each rule set 28 has a unique rule set identifier (UUID), rule set name, rule set
24 description, rule set version, rule set type (e.g. controlled by system 10 or by user), and a rule
25 set category (e.g. generic rule set categorization such as business, configuration etc.).

26 **[00134]** As described above, each rule set 28 is also comprised of one or more rules.

27 **Rule Definition**

[00135] A rule is conceptually a form of metadata that specifies a parameter, tests to detect an incompatibility between the baseline and target systems, 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
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. 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 12, 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.

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

Rule Set Example

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

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

[00150] Rule 6 scrutinizes the lists of all patches applied to the source and target systems and penalizes cases where they differ. The mutex flag is set, indicating that the penalty is applied only once, no matter how many patch differences exist. This rule is ignored in cases where either the operating system or operating system version are different. Rule 7 penalizes system combinations of servers that are running the same OS but are configured to run a different number of kernel bits (e.g. 64-bit vs 32-bit). Rule 8 penalizes combinations where

1 there are kernel parameters defined on the source that are not defined on the target. This rule
2 is not applied if the operating systems are different.

3 **[00151]** Rule 9 scrutinizes a specific kernel setting (SHMMAX, the setting that specifies how
4 much shared memory a system can have) and penalizes combinations where it is set to a lower
5 value on the target than it is on the source system. Rule 10 penalizes combinations of systems
6 that are running different database version, e.g. Oracle™ 9 vs. Oracle™ 8. Rule 11 penalizes
7 combinations of systems that are running different versions of Oracle™. Rule 11 is suppressed
8 if the more specific Rule 10 is true. It should be noted that the remediation cost is relatively
9 high, owing to the fact that it will take a software upgrade to eliminate this discrepancy. In some
10 cases the remediation cost can be low where the upgrade is less expensive. Rule 12 penalizes
11 combinations of systems that are running different versions of Apache. It should be noted that
12 the remediation cost is relatively low, as apache is an open source product and the cost of
13 upgrade is based on the hourly cost of a system administrator and how long it will take to
14 perform the upgrade.

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

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

[00154] It will be appreciated that the above described rules and rule set 28 are shown for 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 16 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 16 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
-------	-------------

Field	Description
Name	As before
Description	As before
Data Query Type	As before
Query Value	As before
Test	Test expression
Weight	As before
Mutex Flag	As before
Match Flag	As before
Match Test	Rule dependency expression to determine whether to perform the corresponding match action
Match Action	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
7 stop.

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

16 **1-to-1 Rule-based Compatibility Analysis**

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

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

1 item. As part of the rule evaluation and subsequent compatibility score calculation, if a rule is
2 true, the corresponding cost to address the deficiency is added to the remediation cost for the
3 pair of systems 16 being analysed.

4 **[00175]** The evaluation of the rules is shown in Figure 18. The evaluation of the rules
5 considers the snapshot data 18 for the source system and the target system, as well as the
6 differential rule set 28 that being applied. For each rule in the set 28, the data referenced by the
7 rule is obtained for both the target and source. The rule is evaluated by having the rule engine
8 90 compare the data. If the rule is not true (i.e. if the systems 16 are the compatible according
9 to the rule definition) then the data 18 is not considered in the compatibility score and the next
10 rule is evaluated. If the rule is true, the rule details are added to an intermediate result. The
11 intermediate result includes all true rules.

12 **[00176]** Preferably, a suppression tag is included with each rule. As discussed above, the
13 suppression tag indicates other rules that are not relevant if that rule is true. The suppression
14 flag allows the program 10 to avoid unnecessary computations. A mutex flag is also preferably
15 used to avoid unfairly reducing the score for each true rule when the rules are closely affected
16 by each other.

17 **[00177]** Once each rule has been evaluated, a list of matched rules is created by removing
18 suppressed rule entries from the intermediate results based on rule dependencies, which are
19 defined by rule matching and suppression settings (e.g. match flags and suppression tags).
20 The compatibility score for that particular source/target pair is then computed based on the
21 matched rules, weights and mutex settings. Remediation costs are also calculated based on
22 the cost of updating/upgrading etc. and the mutex settings.

23 **[00178]** Turning back to Figure 17, the current target is then evaluated against all remaining
24 sources and then the next target is evaluated. As a result, an N x N map 32 can be created that
25 shows a compatibility score for each system against each other system. The map 32 can be
26 sorted by grouping the most compatible systems. The sorted map 32 is comprised of every
27 source/target combination and thus provides an organized view of the compatibilities of the
28 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 Type	Benchmark type corresponding to workload type	cpu
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

Property	Description	Examples
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 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 system (optional)	cpuBenchmark
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 16 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
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

Timestamp	%CPU used (A)	%CPU used (B)	%CPU used (C)
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, 1st quartile, median, 3rd 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. 3rd quartile) while reducing sensitivity to outlier values.

Workload Data Extraction

[00192] Workload data is typically collected and stored in the workload data cache for each system 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 Nth percentile (e.g. 95th) 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 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:

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

$$\text{3rd Quartile} = \text{Q3}_A + \text{Q3}_B + \text{Q3}_C = 110\%$$

1 **[00198]** Median = $Q2_A + Q2_B + Q2_C$ = 70%

$$2 \quad [00199] \quad 1^{\text{st}} \text{ Quartile} \quad = \quad Q1_A + Q1_B + Q1_C \quad = \quad 35\%$$

3 **[00200]** Minimum = Min_A + Min_B + Min_C = 10%

4 **[00201]** The resulting sums can exceed theoretical maximum capacity/benchmark values
5 (e.g. >100% CPU used).

6 **[00202]** The cascade strategy processes the statistical values in descending order, starting
7 with the highest statistical value (i.e. maximum value). The strategy adds like statistical values
8 as with original, but may clip the resulting sums if they exceed a configurable limit and cascades
9 a portion of the excess value to the *next statistic* (i.e. the excess of sum of the maximum values
10 is cascaded to 3rd quartile). The relevant configuration properties for the cascade calculation
11 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
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

12 *Table 9: Workload Cascade Configuration*

13 **[00203]** The following example applies cascade calculation to the example workload data
14 from Table 2 with the following settings:

```
15  [00204] wci.cascade_overflow = true
```

```
16  [00205] wci.cascade_overflow_proportion = 0.5
```

```
17 [00206] wci.clip_type = max
```

[00207] 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:

$$[00208] \quad \text{Max}_{\text{Original}} = \text{Max}_A + \text{Max}_B + \text{Max}_C$$

$$[00209] \quad = 180\%$$

$$[00210] \quad \text{Max}_{\text{Clipped}} = \text{Minimum of } (\text{Max}_{\text{Original}}, \text{clipping level})$$

$$[00211] \quad = \text{Minimum of } (180, 100)$$

$$[00212] \quad = 100\%$$

$$[00213] \quad \text{Max}_{\text{Excess}} = \text{Max}_{\text{Original}} - \text{Max}_{\text{Clipped}}$$

$$[00214] \quad = 180\% - 100\%$$

$$[00215] \quad = 80\%$$

$$[00216] \quad \text{Max}_{\text{Overflow}} = \text{Max}_{\text{Excess}} * \text{wci.cascade_overflow_proportion}$$

$$[00217] \quad = 80\% * 0.5$$

$$[00218] \quad = 40\%$$

$$[00219] \quad \text{Q3}_{\text{Original}} = \text{Q3}_A + \text{Q3}_B + \text{Q3}_C$$

$$[00220] \quad = 110\%$$

$$[00221] \quad \text{Q3}_{\text{Cascade}} = \text{Q3}_{\text{Original}} + \text{Max}_{\text{Overflow}}$$

$$[00222] \quad = 110\% + 40\%$$

$$[00223] \quad = 150\%$$

1	[00224]	$Q3_{\text{Clipped}}$	=	Minimum of ($Q3_{\text{Cascade}}$, clipping level)
2	[00225]		=	Minimum of (150, 100)
3	[00226]		=	100%
4	[00227]	$Q3_{\text{Excess}}$	=	50%
5	[00228]	$Q3_{\text{Overflow}}$	=	25%
6	[00229]	$Q2_{\text{Original}}$	=	70%
7	[00230]	$Q2_{\text{Cascade}}$	=	$Q2_{\text{Original}} + Q3_{\text{Overflow}}$
8	[00231]		=	70% + 25%
9	[00232]		=	95%
10	[00233]	$Q2_{\text{Clipped}}$	=	Minimum of ($Q2_{\text{Cascade}}$, clip level)
11	[00234]		=	Minimum of (95, 100)
12	[00235]		=	95%
13	[00236]	$Q2_{\text{Overflow}}$	=	0%
14	[00237]	$Q1_{\text{Original}}$	=	35%
15	[00238]	$Q1_{\text{Cascade}}$	=	35%
16	[00239]	$Q1_{\text{Clipped}}$	=	35%
17	[00240]	$Q1_{\text{Overflow}}$	=	0%
18	[00241]	$\text{Min}_{\text{Original}}$	=	10%
19	[00242]	$\text{Min}_{\text{Clipped}}$	=	10%

1 **[00243]** The consolidated statistics for the above example are summarized in the following
 2 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	100	80	40
Q3	110	150	100	50	25
Q2	70	95	95	0	0
Q1	35	35	35	0	0
Min	10	10	10	0	0

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

4 **[00244]** Similarly, the following example applies cascade calculation to the example workload
 5 data from Table 7 with the following settings:

6 **[00245]** wci.cascade_overflow = true

7 **[00246]** wci.cascade_overflow_proportion = 0.5

8 **[00247]** wci.clip_type = limit

9 **[00248]** wci.clip_limit_ratio = 2

10 **[00249]** This example specifies a *limit* clip type to indicate that the clipping level is based on
 11 the analysis threshold for the workload type. The *clip_limit_ratio* specifies the ratio to apply to
 12 the threshold to calculate the actual clipping level. For instance, if the threshold is 80% and the
 13 clip limit ratio is 2, the clipping level is 160%. The consolidated workload statistics based on the
 14 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	10	0	0
-----	----	----	----	---	---

Table 11: Cascaded Statistics (Example 2)

Workload Compatibility Scoring

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

[00251] Workloads are assessed separately for two scenarios: like-times and worst case. The like times scenario combines the workload of the systems at like times (i.e. same hours) for the 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, 3rd quartile, median, 1st 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 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,

[00260] Penalty = 1

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

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

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

[00264] Penalty_{Max} = (160 – 80) / (160 – 80)

[00265] = 1

[00266] Penalty_{Q3} = (120 – 80) / (160 – 80)

[00267] = 0.5

[00268] Penalty_{Q2} = 0 [since 70 < 80]

[00269] Penalty_{Q1} = 0 [since 35 < 80]

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

[00271] The workload score is composed of the weighted penalty values. Penalty weights can 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 rd Quartile	Like Times	wci.score.q3_like_times	0.4
3 rd 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 st Quartile	Like Times	wci.score.q1_like_times	0
1 st 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

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

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

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

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

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

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

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

2 **[00282]** $= 30$

3 **1-to-1 Workload Compatibility Analysis Process Flow**

4 **[00283]** A flow chart illustrating a workload compatibility analysis is shown in Figure 21.

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

8 **[00284]** The workload analysis considers one or more workload types, e.g. CPU busy, the
9 workload limits 94, e.g. 75% of the CPU being busy, and the system benchmarks 96, e.g.
10 relative CPU power. Each system 16 in the workload data set is considered as a target (T = 1
11 to N) and compared to each other system 16 in the data set 18 as the source (S = 1 to N). The
12 analysis engine 64 first determines if the source and target are the same. If yes, then the
13 workload compatibility score is set to 100 and no additional analysis is required for that pair. If
14 the source and target are different, the system benchmarks are then used to normalize the
15 workloads (if required). The normalized source workload histogram is then stacked on the
16 normalized target system.

17 **[00285]** System benchmarks can normalize workloads as follows. For systems X and Y, with
18 CPU benchmarks of 200 and 400 respectively (i.e. Y is 2x more powerful than X), if systems X
19 and Y have average CPU utilization of 10% and 15% respectively, the workloads can be
20 normalized through the benchmarks as follows. To normalize X's workload to Y, multiply X's
21 workload by the benchmark ratio X/Y, i.e. $10\% \times 200/400 = 5\%$. Stacking X onto Y would then
22 yield a total workload of $5\% + 15\% = 20\%$. Conversely, stacking Y onto X would yield the
23 following total workload: $10\% + 15\% \times 400/200 = 40\%$.

24 **[00286]** Using the stacked workload data, the workload compatibility score is then computed
25 for each workload type as described above.

26 **[00287]** Each source is evaluated against the target, and each target is evaluated to produce
27 an N x N map 34 of scores, which can be sorted to group compatible systems (see Figure 5).

1 Preferably, a workload compatibility results is generated that includes the map 34 and workload
2 compatibility scoring details and normalized stacked workload histograms that can be viewed by
3 selecting the appropriate cell 92 (see Figure 22). The workload compatibility results are then
4 combined with the rule-based compatibility results to produce the overall compatibility scores,
5 described below.

6 **1-to-1 Overall Compatibility Score Calculation**

7 **[00288]** The results of the rule and workload compatibility analyses are combined to compute
8 an overall compatibility score for each server pair. These scores preferably range from 0 to
9 100, where higher scores indicate greater compatibility and 100 indicating complete or 100%
10 compatibility.

11 **[00289]** As noted above, the analysis input can include importance factors. For each rule set
12 28 and workload type 30 included in the analysis, an importance factor 88 can be specified to
13 adjust the relative contribution of the corresponding score to the overall score. The importance
14 factor 88 is an integer, preferably ranging from 0 to 10. A value of 5 has a neutral effect on the
15 contribution of the component score to the overall score. A value greater than 5 increase the
16 importance whereas a value less than 5 decreases the contribution.

17 **[00290]** The overall compatibility score for the system pair is computed by combining the
18 individual compatibility scores using a formula specified by an overlay algorithm which performs
19 a mathematical operation such as multiply or average, and the score is recorded.

20 **[00291]** Given the individual rule and workload compatibility scores, the overall compatibility
21 score can be calculated by using the importance factors as follows for a “multiply” overlay:

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

23 **[00292]** where O is the overall compatibility score, n is the total number of rule sets 28 and
24 workload types 30 included in the analysis, S_i is the compatibility score of the i^{th} rule set 28 or
25 workload type 30 and F_i is the importance factor of the i^{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 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.

1 **[00300]** It is therefore seen that the graphical output of the map 32 provides an intuitive
2 mapping between the source/target pairs in the environment 12 to assist in visualizing where
3 compatibilities exist and do not exist. In Figure 4, it can be seen that a system pair having a
4 score = 100 indicates complete compatibility between the two systems 16 for the particular
5 strategy being observed, e.g. based on a chosen rule set(s) 28. It can also be seen that a
6 system pair with a relatively lower score such as 26 is relatively less compatible for the strategy
7 being observed.

8 **[00301]** The detailed differences shown in Figures 19 can be viewed by clicking on the
9 relevant cell 92. Selecting a particular cell 92 accesses the detailed differences table 100
10 shown in Figure 19 which shows the important differences between the two systems, the rules
11 and weights that were applied and preferably a remediation cost for making the servers more
12 compatible. As shown in Figure 19, a summary differences table 102 may also be presented
13 when selecting a particular cell 92, which lists the description of the differences and the weight
14 applied for each difference, to give a high level overview of where the differences arise.

15 **System Workload Compatibility Visualization**

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

1 **[00303]** The workload analysis details shown in Figures 22 can be viewed by clicking on the
2 relevant cell 92. Selecting a particular cell 92 accesses the information about the workload
3 analysis that generated the score shown in Figure 22, which shows the key stacked workload
4 values, the workload benchmarks that were applied and preferably workload charts for each
5 system separately, and stacked together.

6 7 **Overall Compatibility Visualization**

8 **[00304]** An example overall compatibility analysis map 36 is shown in Figure 6. The map 36
9 comprises a similar arrangement as the maps 32 and 34, which lists the servers in the
10 uppermost row and leftmost column to provide 100% compatibility along the diagonal.
11 Preferably the same scoring and shading convention is used by all types of compatibility maps.
12 The map 36 provides a visual display of scoring for candidate system pairs that considers the
13 rule-based compatibility maps 32 and the workload compatibility maps 34.

14 **[00305]** The score provided in each cell 92 indicates the overall compatibility for
15 consolidating systems 16. It should be noted that in some cases two systems 16 can have a
16 high configuration compatibility but a low workload compatibility and thus end up with a reduced
17 or relatively low overall score. It is therefore seen that the map 36 provides a comprehensive
18 score that considers not only the compatibility of systems 28 at the setting level but also in its
19 utilization. By displaying the configuration maps 32, business maps, workload maps 34 and
20 overall map 36 in a consolidation roadmap, a complete picture of the entire system can be
21 ascertained in an organized manner. The maps 32, 34 and 36 provide a visual representation
22 of the compatibilities and provide an intuitive way to evaluate the likelihood that systems can be
23 consolidated, to analyse compliance and drive remediation measures to modify systems 16 so
24 that they can become more compatible with other systems 16 in the environment 12. It can
25 therefore be seen that a significant amount of quantitative data can be analysed in a convenient
26 manner using the graphical maps 32, 34 and 36, and associated reports and graphs (described
27 below).

28 **[00306]** For example, a system pair that is not compatible only for the reason that certain
29 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.

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

Sorting Examples

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

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

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

[00311] Sorting by group involves computing the difference between each system pair, and arranging the systems to minimize the total difference between each pair of adjacent systems in the map. The difference between a system pair can be computed by taking the square root of the sum of the squares of the difference of a pair's individual compatibility score against each other system in the analysis. In general, the smaller the total difference between two systems, the more similar the two systems with respect to their compatibility with the other systems. The group sort promotes the visualization of the logical breakdown of an environment by producing clusters of compatible systems 18 around the map diagonal. These clusters are indicative of

compatible regions in the environment 12. In virtualization analysis, these are often referred to as “affinity regions.”

Analysis Results-User Interaction

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

Multi-Dimensional Compatibility Analysis

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

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

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

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

- Systems: S1, S2, S3, ... S20
- Analyzed with rule sets: R1, R2, R3
- Analyzed with workload types: W1, W2
- Transfer sets:
 - T1 (S1, S2, S3 stacked onto S4)
 - T2 (S5, S6, S7, S8, S9 stacked onto S10)
 - T3 (S11, S12, S13 stacked onto S14)
 - T4 (S15 stacked onto S16)
 - T5 (S17 stacked onto S18)
- Unaffected systems: S19, S20

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

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

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

[00320] Similar to the 1-to-1 compatibility analysis, this analysis involves 4 stages. The first stage is gets the system data required for the analysis to produce the analysis data snapshot. The second stage performs a multi-dimensional compatibility analysis for each rule set for each transfer set. Next, the workload compatibility analysis is performed for each workload type for each transfer set. Finally, these analysis results are combined to determine overall 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
10 and 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
16 or 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-
28 compatibility between the source entities and the target as the primary analysis mechanism, but,

depending on the rule sets in use, may also consider the N-by-N intra-compatibility of the source entities amongst each other.

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

N-to-1 Intercompatibility Score Calculation

[00327] N-to-1 intercompatibility scores reflect the compatibility between N source entities and a single target as defined by a transfer set 23 as shown in Figure 24(c). This analysis is performed with respect to a given rule set and involves: 1) Separately evaluate each source entity against the target 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:

$$\text{[00328]} \quad S = 100 * (1 - w_1) * (1 - w_2) * (1 - w_3) * \dots * (1 - w_n);$$

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

N-to-1 Score Example

[00330] 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 s16, s2 against s16 and s3 against s16. Assuming this analysis produces a list of matched items comprising those shown below in Table 14:

#	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

[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	0.03
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	0.02
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	0.01
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 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.

1 **[00340]** For each workload type 30, each transfer set 23 is evaluated. For each source in the
2 transfer set 23, the system benchmarks are used to normalize the workloads as discussed
3 above, and the source workload is stacked on the target system. Once every source in the set
4 is stacked on the target system, the workload compatibility score is computed as discussed
5 above. The above is repeated for each transfer set 23. A multi-stack report may then be
6 generated, which gives a workload compatibility scorecard for the transfer sets along with
7 workload compatibility scoring details and normalized multi-stacked workload charts.

8 **[00341]** Sample multi-dimensional compatibility analysis results are shown in Figure 26-28.
9 Figure 26 shows a compatibility score map 110 with the analyzed transfer sets. Figures 27 and
10 28 show the summary 112 and charts 114 from the multi-stack workload compatibility analysis
11 details.

12 **Consolidation Analysis**

13 **[00342]** The consolidation analysis process flow is illustrated as D in Figure 9. Using the
14 common compatibility analysis input and additional auto fit inputs, this analysis seeks the
15 consolidation solution that maximizes the number of transfers while still fulfilling the several pre-
16 defined constraints. The consolidation analysis repeatedly employs the multi-dimensional
17 compatibility analysis to assess potential transfer set candidates. The result of the consolidation
18 analysis comprises of the consolidation solution and the corresponding multi-dimensional
19 compatibility analysis.

20 **[00343]** A process flow of the consolidation analysis is shown in Figure 29.

21 **[00344]** The auto fit input includes the following parameters: transfer type (e.g. virtualize or
22 stacking), minimum allowable overall compatibility score for proposed transfer sets, minimum
23 number of source entities to transfer per target, maximum number of source entities to transfer
24 per target, and quick vs. detailed search for the best fit. Target systems can also be designated
25 as malleable or concrete models.

26 **[00345]** As part of a compatibility analysis input specification, systems can be designated for
27 consideration as a source only, as a target only or as either a source or a target. These
28 designations serve as constraints when defining transfers in the context of a compatibility

1 analysis. The analysis can be performed on an analysis with pre-existing source-target
2 transfers. Analyses containing systems designated as source or target-only (and no source or
3 target designations) are referred to as “directed analysis.”

4 **[00346]** The same transfer type may be assumed for all automatically determined transfers
5 within an analysis. The selected transfer type affects how the compatibility analysis is
6 performed. The minimum overall compatibility score dictates the lowest allowable score
7 (sensitivity) for the transfer sets to be included in the consolidation solution. Lowering the
8 minimum allowable score permits a greater degree of consolidation and potentially more
9 transfers. The minimum and maximum limits for source entities to be transferred per target
10 (cardinality) define additional constraints on the consolidation solution. The quick search
11 performs a simplified form of the auto fit calculation, whereas the detailed search performs a
12 more exhaustive search for the optimal solution. This distinction is provided for quick
13 assessments of analyses containing a large numbers of systems to be analyzed.

14 **[00347]** The transfer auto fit problem can be considered as a significantly more complex form
15 of the classic bin packing problem. The bin packing problem involves packing objects of
16 different volumes into a finite number of bins of varying volumes in a way that minimizes the
17 number of bins used. The transfer auto fit problem involves transferring source entities onto a
18 finite number of targets in a way that maximizes the number of transfers. The basis by which
19 source entities are assessed to “fit” onto targets is based on the highly nonlinear compatibility
20 scores of the transfer sets. As a further consideration, which can increase complexity, some
21 entities may be either source or targets. The auto fit problem is a combinatorial optimization
22 problem that is computationally expensive to solve through a brute force search of all possible
23 transfer set permutations. Although straightforward to implement, this exhaustive algorithm is
24 impractical due to its excessive computational and resource requirements for medium to large
25 data sets. Consequently, this class of problem is most efficiently solved through heuristic
26 algorithms that yield good but likely suboptimal solutions.

27 **[00348]** There are four variants of the heuristic auto fit algorithm that searches for the best
28 consolidation solution:

29 **[00349]** Quick Stack – quick search for a stacking-based consolidation solution;

1 **[00350]** Detailed Stack – more comprehensive search for a stacking-based consolidation
2 solution;

3 **[00351]** Quick Virtualization – quick search for a virtualization-based consolidation solution;
4 and

5 **[00352]** Detailed Virtualization – more comprehensive search for a virtualization-based
6 consolidation solution.

7 **[00353]** The auto fit algorithms are iterative and involve the following common phases:

8 **Compile Valid Source and Target Candidates**
9

10 **[00354]** The initial phase filters the source and target lists by eliminating invalid entity
11 combinations based on the 1-to-1 compatibility scores that are less than the minimum allowable
12 compatibility score. It also filters out entity combinations based on the source-only or target-
13 only designations.

14 **Set up Auto Fit Parameters**

15 **[00355]** The auto fit algorithm search parameters are then set up. The parameters can vary
16 for each algorithm. Example search parameters include the order by which sources and targets
17 are processed and the criteria for choosing the best transfer set 23.

18 **Compile Candidate Transfer Sets**
19

20 **[00356]** The next phase compiles a collection of candidate transfer sets 23 from the available
21 pool of sources and targets. The candidate transfer sets 23 fulfill the auto fit constraints (e.g.
22 minimum allowable score, minimum transfers per transfer set, maximum transfers per transfer
23 set). The collection of candidate transfer sets may not represent a consolidation solution (i.e.
24 referenced sources and targets may not be mutually exclusive amongst transfer sets 23). The
25 algorithms vary in the criteria employed in composing the transfer sets. In general, the detailed
26 search algorithms generate more candidate transfer sets than quick searches in order to assess
27 more transfer permutations.

28 **Choose Best Candidate Transfer Set**
29

1 **[00357]** The next phase compares the candidate transfer sets 23 and chooses the “best”
2 transfer set 23 amongst the candidates. The criteria employed to select the best transfer set 23
3 varies amongst the algorithms. Possible criteria include the number of transfers, the
4 compatibility score, general compatibility of entities referenced by set and whether the transfer
5 set target is a target-only.

6 **Add Transfer Set to Consolidation Solution**

7 **[00358]** Once a transfer set is chosen, it is added to the intermediate consolidation solution.
8 The entities referenced by the transfer set are removed from the list of available sources and
9 targets and the three preceding phases are repeated until the available sources or targets are
10 consumed.

11 **Compile Consolidation Solution Candidates**

12
13 **[00359]** Once all the sources or targets are consumed or ruled out, the consolidation solution
14 is considered complete and added to a list of candidate solutions. Additional consolidation
15 solutions can be compiled by iterating from the second phase with variations to the auto fit
16 parameters for compiling and choosing candidate transfer sets.

17 **Choose Best Consolidation Solution**

18
19 **[00360]** The criteria used to stop compiling additional solutions can be based on detecting
20 that the solution is converging on a pre-defined maximum number of iterations.

21 **[00361]** Finally, the best candidate consolidation solution can be selected based on some
22 criteria such as the largest reduction of systems with the highest average transfer set scores.
23 The general algorithm is shown in the flow diagram depicted in Figure 30.

24 **Auto Fit Example**

25 **[00362]** The following example demonstrates how a variant of the auto fit algorithm searches
26 for a consolidation solution for a compatibility analysis comprised of 20 systems (S1, S2, S3, ...
27 S20) where 15 systems (S1-15) have been designated to be source-only and 5 (S16-20) to be
28 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 stacking onto target; and 2) When choosing the best transfer set 23, choose set with most transfers and if there is a tie, choose the set 23 with the higher score.

3 – Compile Candidate Transfer Sets

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

[00367] For this quick search algorithm, only a single pass is performed for each target so that only one transfer set candidate is created per target. Other search algorithms can perform

multiple passes per target to assess more transfer permutations. The candidate transfer sets 23 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

Table 17: Candidate Transfer Sets

4 – Choose Best Candidate Transfer Set

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

5 – Add Transfer Set to Consolidation Solution

[00369] The transfer set T3 is added to the consolidation solution, and the entities referenced by this transfer set are removed from the target-source candidates list. The updated target-source 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.

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

13 **Multi-dimensional Compatibility Analysis Example**

14 **[00378]** Continuing with the above example, one or more transfer sets 23 can be defined and
15 the transfer sets 23 evaluated through a multi-dimensional compatibility analysis. The transfer
16 sets 23 can be defined through the user interface, and are signified by one or more arrows that
17 connect the source to the target. The color of the arrow can be used to indicate the transfer
18 type (see Figure 46). Once the transfer sets 23 have been defined, the net effect mode may be
19 selected to run the multi-dimensional analysis. In the resulting compatibility maps, the score in
20 the cells that comprise the transfer set 23 reflects the multi-dimensional compatibility score (see
21 Figure 47).

22 **[00379]** The corresponding overall compatibility details report is shown in Figure 48. Note
23 that there are two sources transferred to the single target.

24 **Consolidation Analysis Example**

25 **[00380]** Again continuing from the above example; a consolidation analysis may be
26 performed to search for a consolidation solution by automating the analysis of the multi-
27 dimensional scenarios. The input screen for the consolidation analysis is shown in Figure 49.
28 Users can specify several parameters including the minimum allowable overall compatibility
29 score and the transfer type. As well, users can choose to keep or remove existing transfer sets
30 before performing the consolidation analysis.

1 **[00381]** Once specified, the consolidation analysis can be executed and the chosen transfer
2 sets 23 are presented in the map as shown in Figure 50. The multi-dimensional compatibility
3 score calculations are also used. Finally, a consolidation summary is provided as shown in
4 Figures 51 and 52. Figure 51 shows that should the proposed transfers be applied, 28 systems
5 would be consolidated to 17 systems, resulting in a 39% reduction of systems. Figure 52 lists
6 the actual transfers proposed by the consolidation analysis.

7 **Commentary**

8 **[00382]** Accordingly, the compatibility and consolidation analyses can be performed on a
9 collection of system to 1) evaluate the 1-to-1 compatibility of every source-target pair, 2)
10 evaluate the multi-dimensional compatibility of specific transfer sets, and 3) to determine the
11 best consolidation solution based on various constraints including the compatibility scores of the
12 transfer sets. Though these analyses share many common elements, they can be performed
13 independently.

14 **[00383]** These analyses are based on collected system data related to their technical
15 configuration, business factors and workloads. Differential rule sets and workload compatibility
16 algorithms are used to evaluate the compatibility of systems. The technical configuration,
17 business and workload related compatibility results are combined to create an overall
18 compatibility assessment. These results are visually represented using color coded scorecard
19 maps.

20 **[00384]** It will be appreciated that although the system and workload analyses are performed
21 in this example to contribute to the overall compatibility analyses, each analysis is suitable to be
22 performed on its own and can be conducted separately for finer analyses. The finer analysis
23 may be performed to focus on the remediation of only configuration settings at one time and
24 spreading workload at another time. As such, each analysis and associated map may be
25 generated on an individual basis without the need to perform the other analyses.

26 **[00385]** It will be appreciated that each analysis and associated map discussed above may
27 instead be used for purposes other than consolidation such as capacity planning, regulatory
28 compliance, change, inventory, optimization, administration etc. and any other purpose where
29 compatibility of systems is useful for analyzing systems 16. It will also be appreciated that the

1 program 10 may also be configured to allow user-entered attributes (e.g. location) that are not
2 available via the auditing process and can factor such attributes into the rules and subsequent
3 analysis.

4 **[00386]** It will further be appreciated that although the examples provided above are in the
5 context of a distributed system of computer servers, the principles and algorithms discusses are
6 applicable to any system having a plurality of sub-systems where the sub-systems perform
7 similar tasks and thus are capable theoretically of being consolidation. For example, a local
8 network having a number of personal computers (PCs) could also benefit from a consolidation
9 analysis.

10 **[00387]** Although the invention has been described with reference to certain specific
11 embodiments, various modifications thereof will be apparent to those skilled in the art without
12 departing from the spirit and scope of the invention as outlined in the claims appended hereto.

Claims:

1. A computer-implemented 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 having been computed using parameters of said plurality of computer systems to indicate a comparable compatibility of respective parameters of one of said plurality of computer systems with respect to another of said plurality of computer systems;

b) using said compatibility scores to determine one or more candidate transfer sets from said plurality of computer systems and, each candidate transfer set indicating one or more source computer systems having compatible parameters that enable the source computer system(s) to be transferred to a target computer system;

c) selecting a desired one of said one or more candidate transfer sets according to a predetermined selection criterion and adding said desired one of said one or more candidate transfer sets to a first candidate consolidation solution;

d) repeating b) and c) by, at each iteration, removing computer systems included in said desired one of said one or more candidate transfer sets and using said compatibility scores corresponding to remaining computer systems to generate at least one additional desired candidate transfer set to be added to said first candidate consolidation solution; and

e) providing said first candidate consolidation solution as a desired consolidation solution when no additional candidate transfer sets can be determined in b).

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
8 data related to their technical configuration, business factors and workloads. Differential rule
9 sets and workload compatibility algorithms are used to evaluate the compatibility of systems.
10 The 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

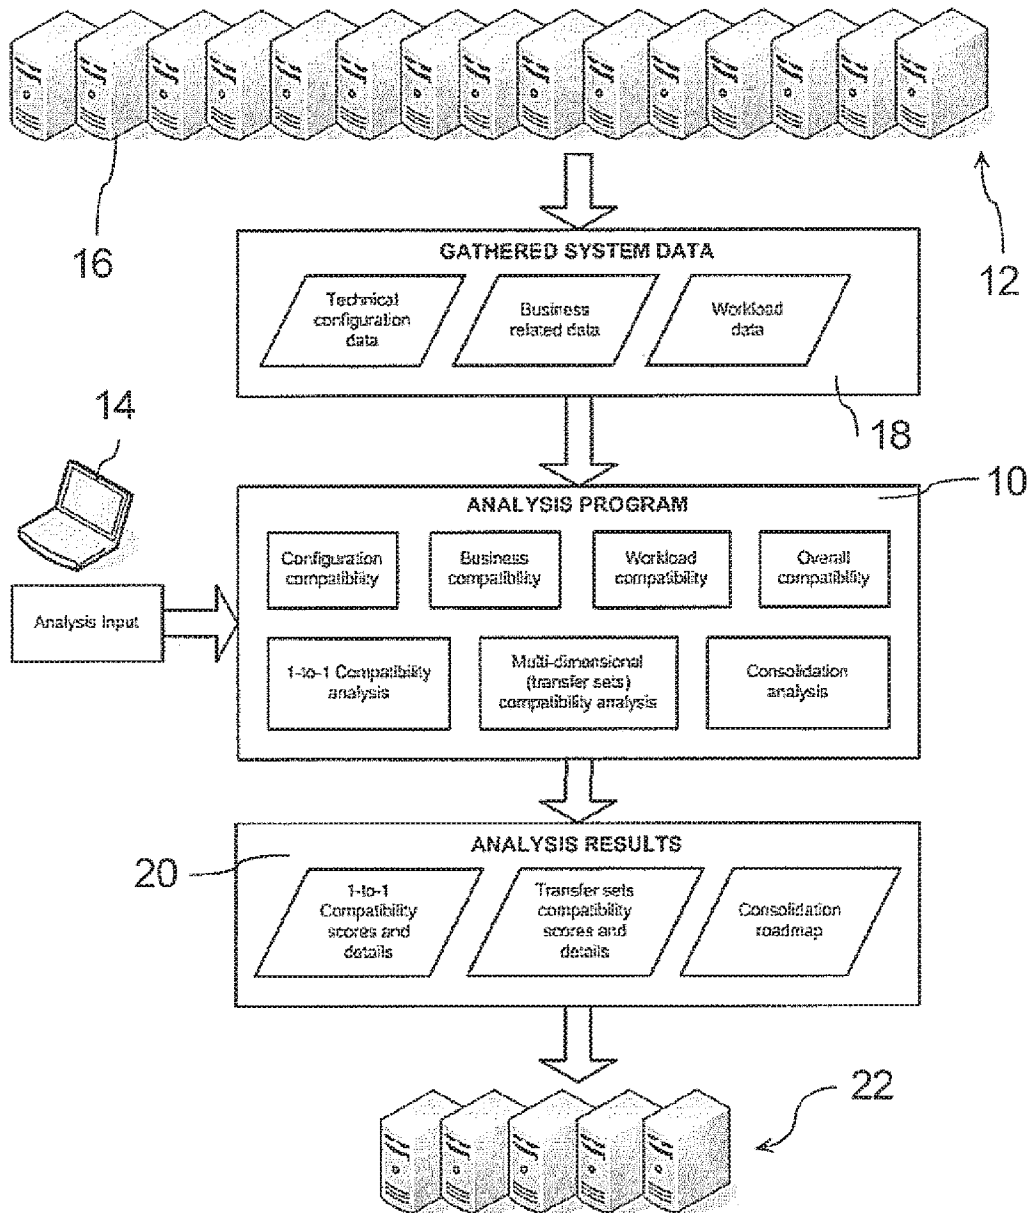


Figure 2




Figure 3

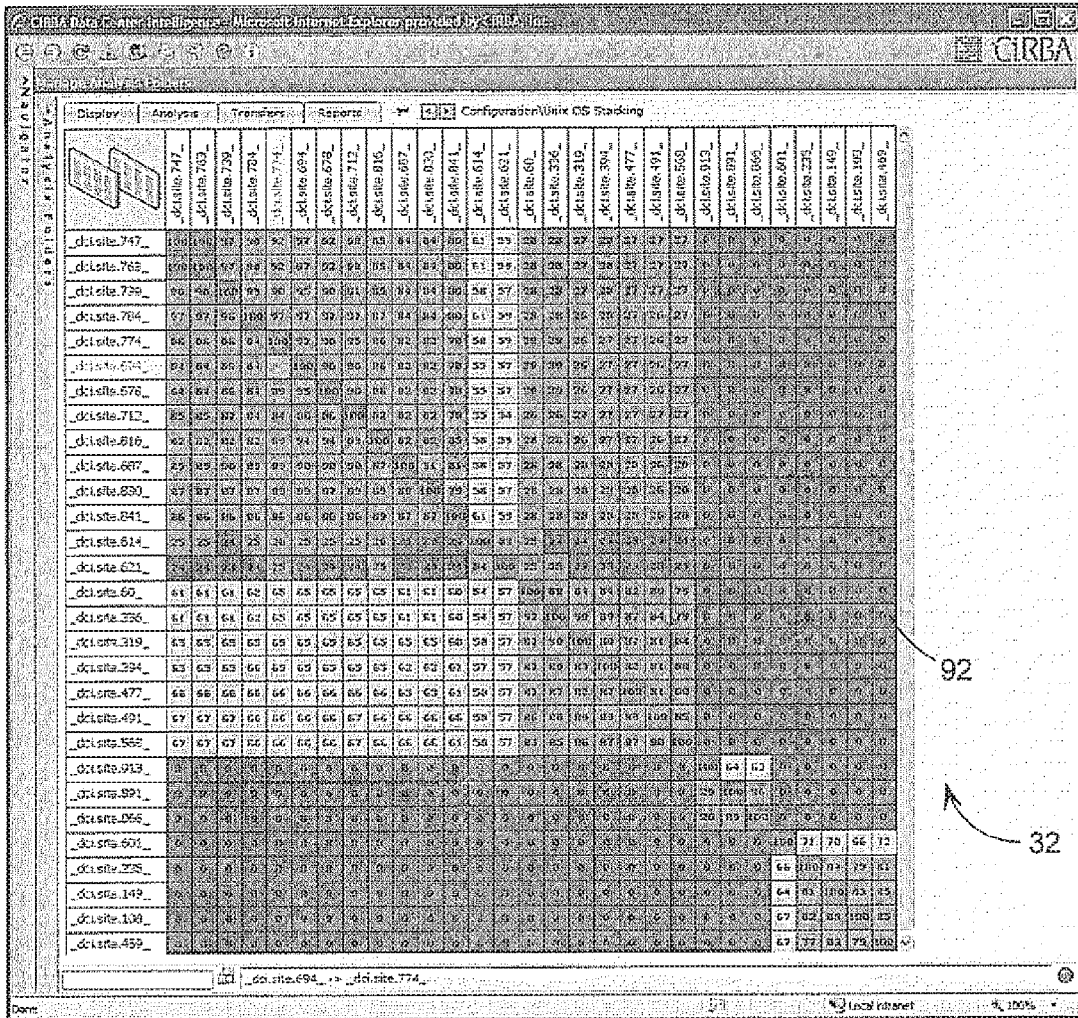


Figure 4

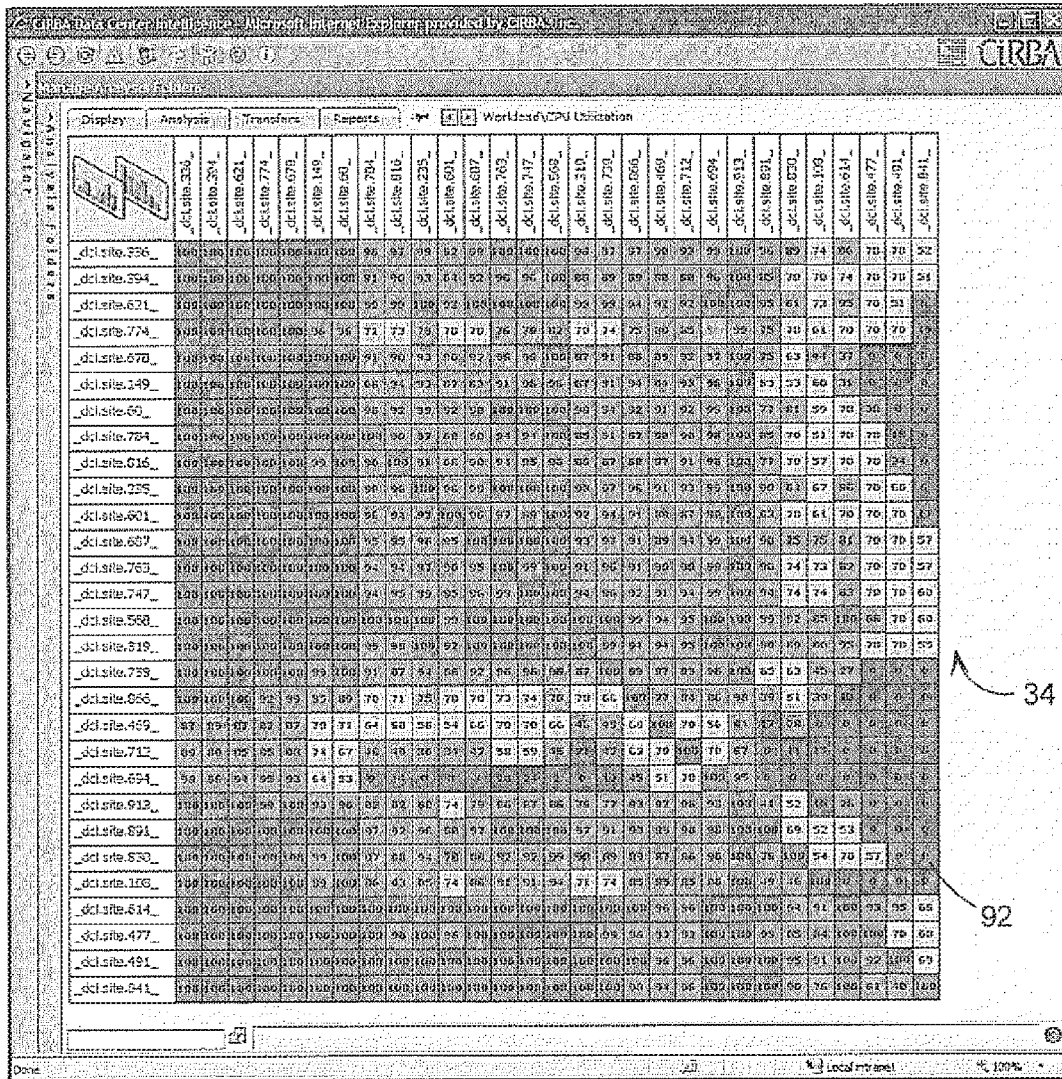


Figure 5

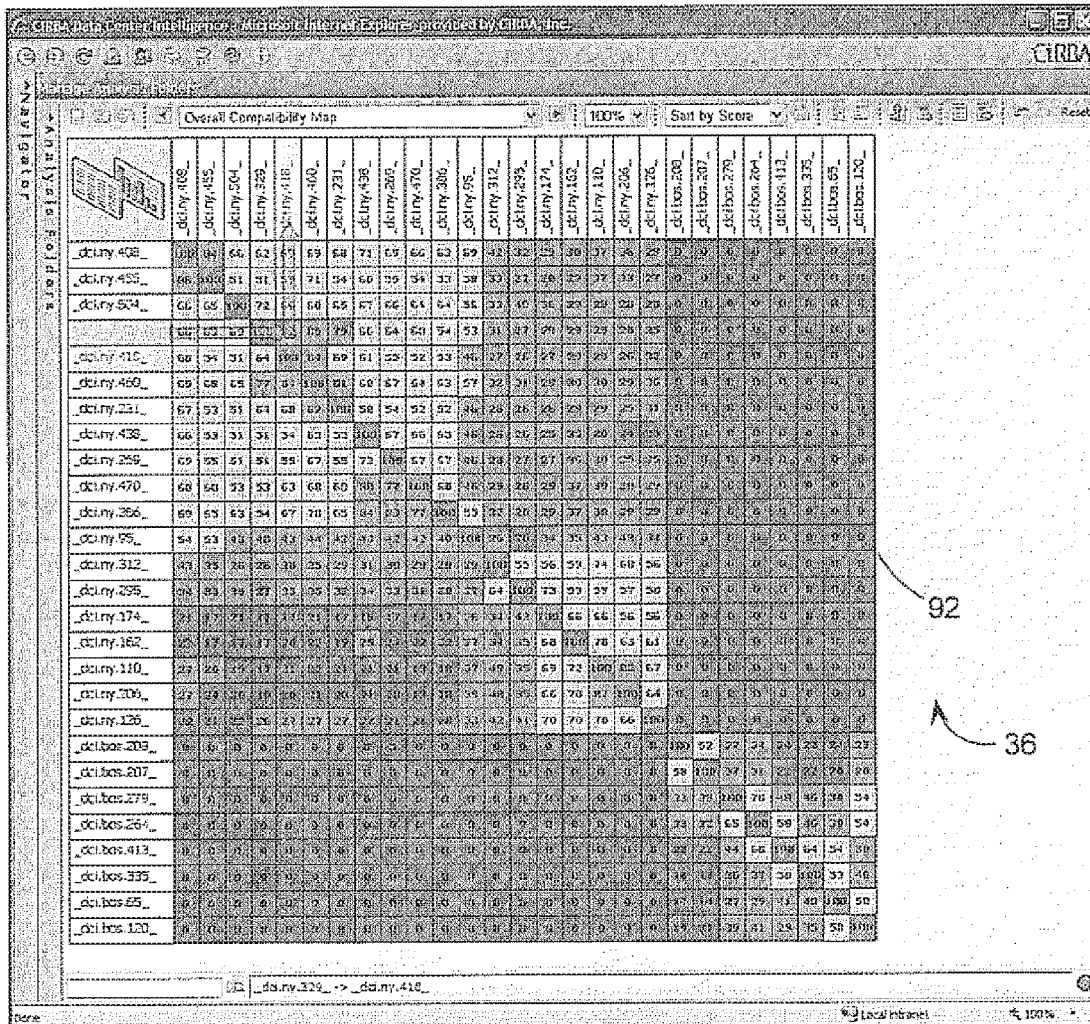


Figure 6




Figure 7

82

Consolidation & Virtualization Analysis Strategy
Instrumentation Enablement

Key:

● Detailed Coverage

○ Partial Coverage or Requires Customization

78

80

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

WebSphere

IIS

Jboss

Apache Tomcat

Virtualization Analysis

UNIX Hardware Variance

Linux Hardware Variance

UNIX Shared OS (Zone) Analysis

Windows Hardware Variance

Windows Shared OS (Zone) Analysis

Configuration Sources (SCI Analysis)

CIRBA Agents	CIRBA Agentless	Third-Party Agents (Direct Audit)	CIRBA Framework Adapters	Data Import
SNMP Agent	SNMP	Compaq IM	Tivoli TME	TADDM
WMI Provider	SSH	Dell OM	VMware ESX	Tivoli CM
	WMI	IBM Director	Concord, etc	
	SQL			
	LMX			

Consolidation & Virtualization Analysis Strategy
Instrumentation Enablement

Key:

- Detailed Coverage
- Partial Coverage or Requires Customization

Figure 8

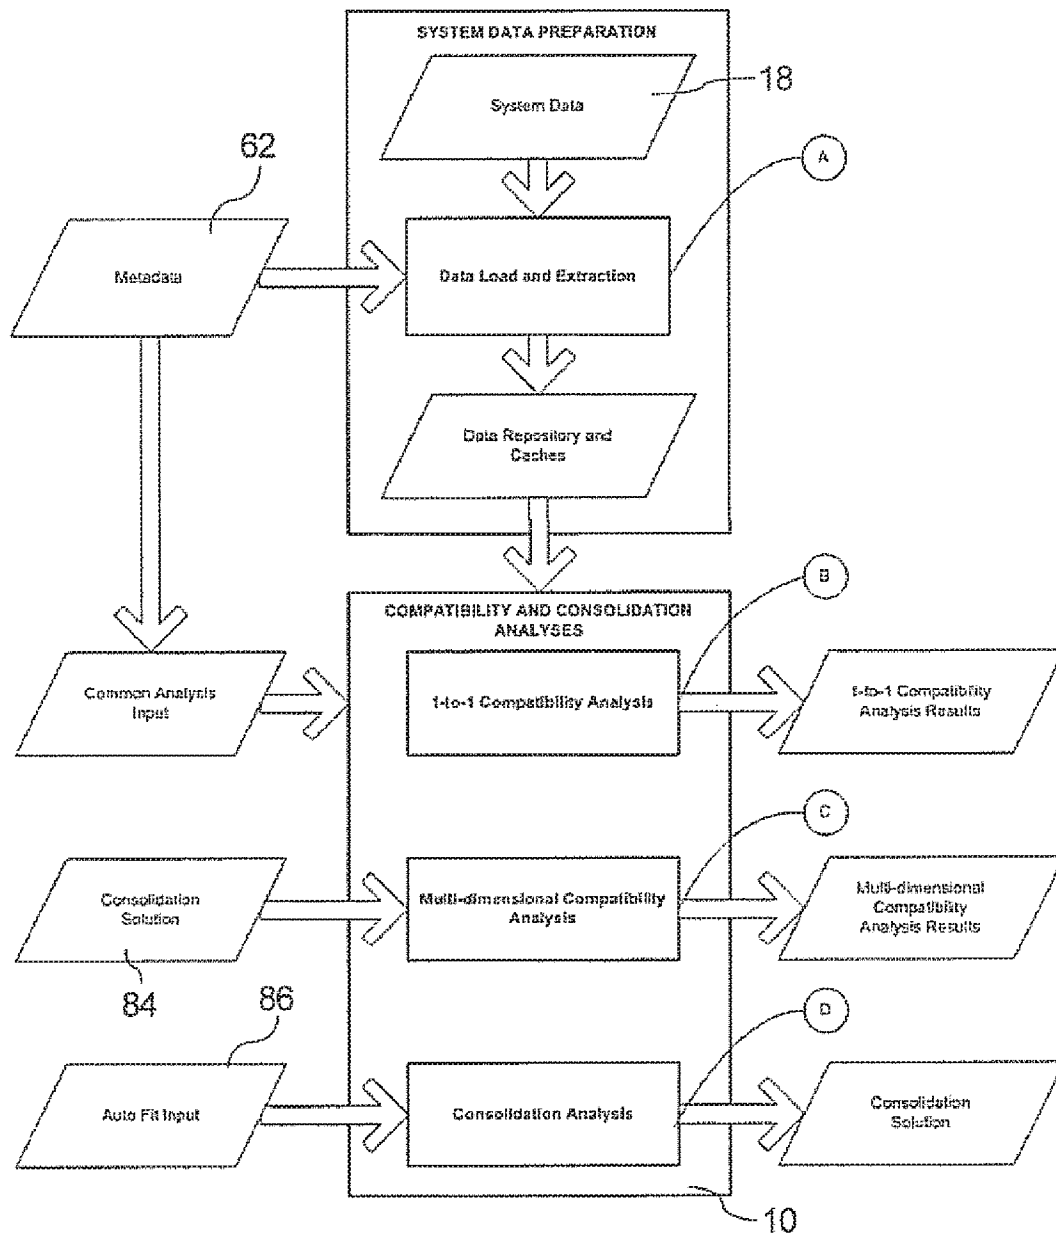


Figure 9

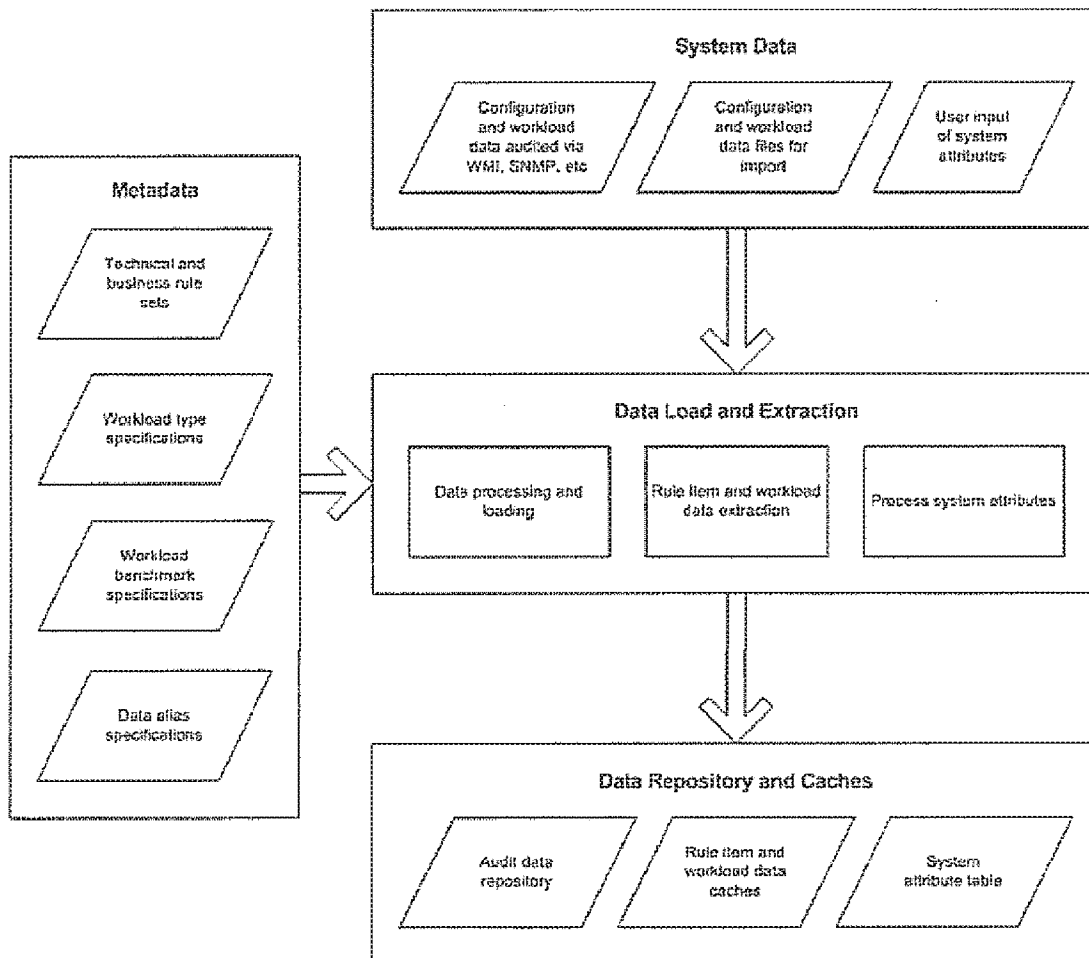


Figure 10

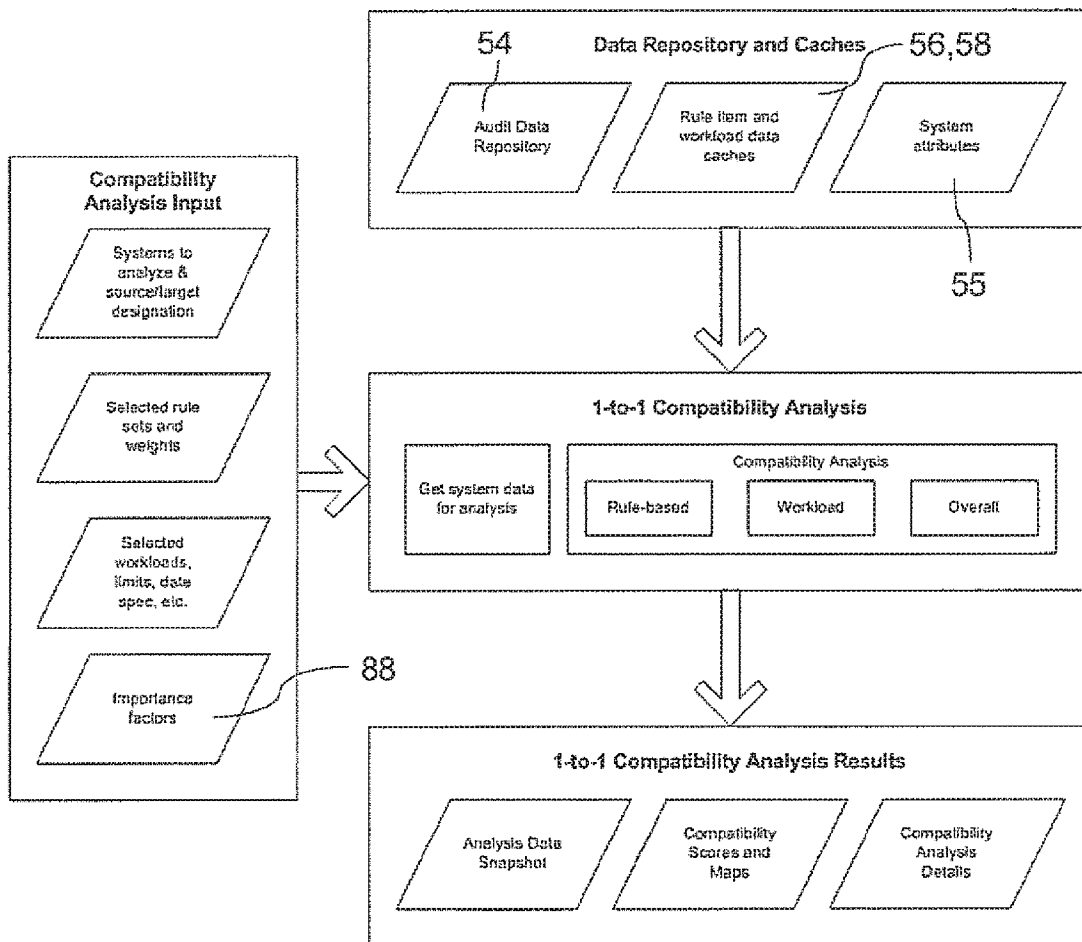


Figure 11

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

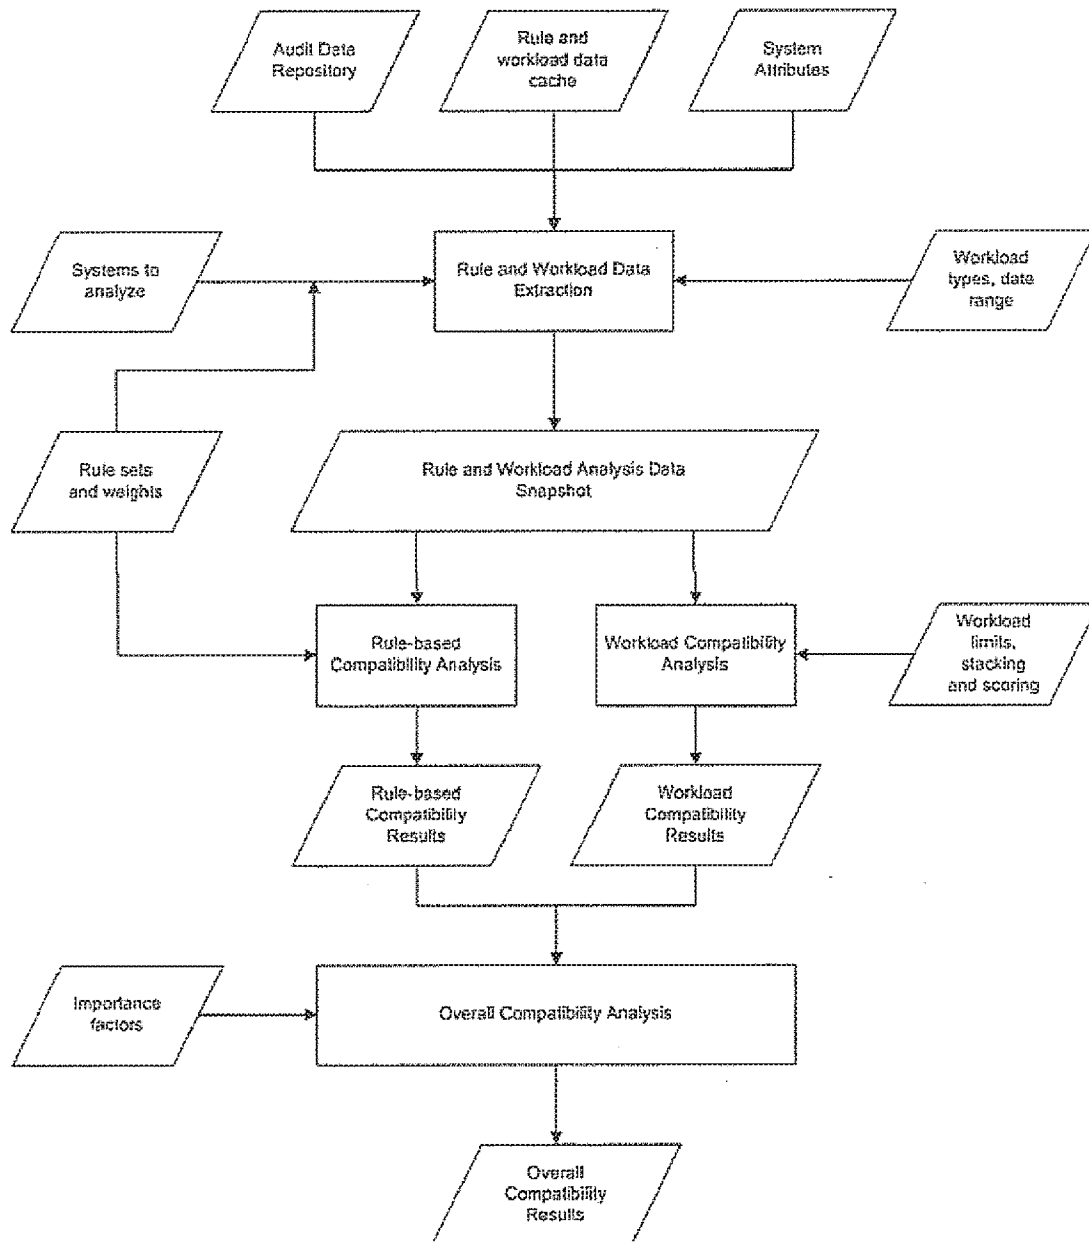


Figure 14

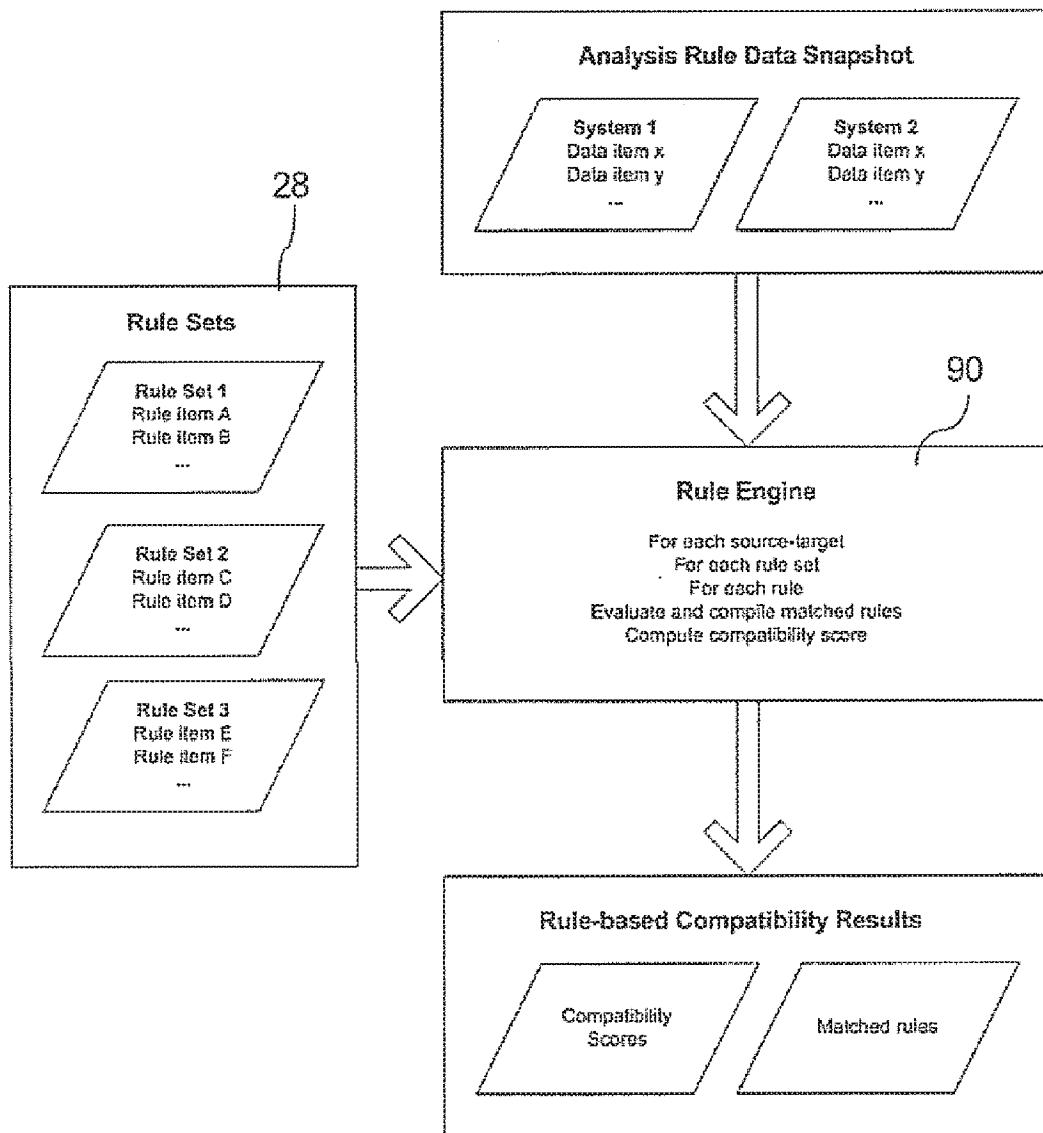


Figure 15

	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 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	\$50	Patch Differences Between Systems
7	UriQuery	cirba-cm/ServerInformation/OSKernelBits			10%			OS	\$50	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/Settling#SHMMAX		<	5%				\$60	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/wln-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	Target
20	UriQuery	win32-application/win-product#McAfee VirusScan Enterprise	present	absent	5%			OS	\$200	McAfee Not On Target

Figure 16

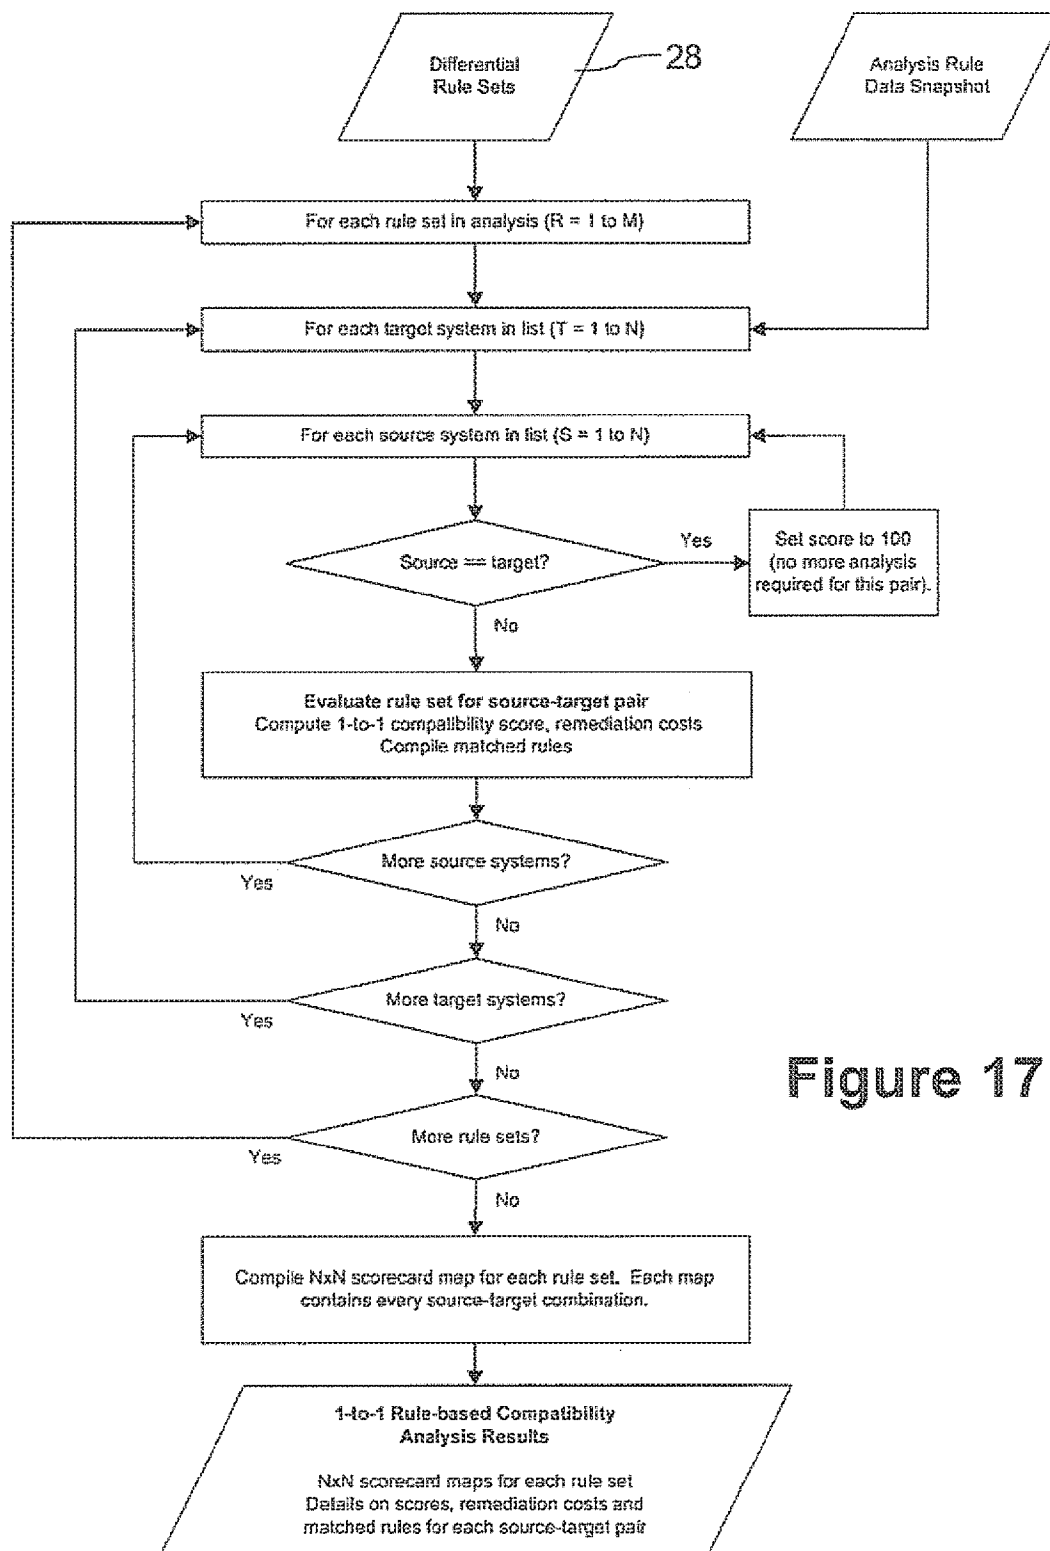


Figure 17

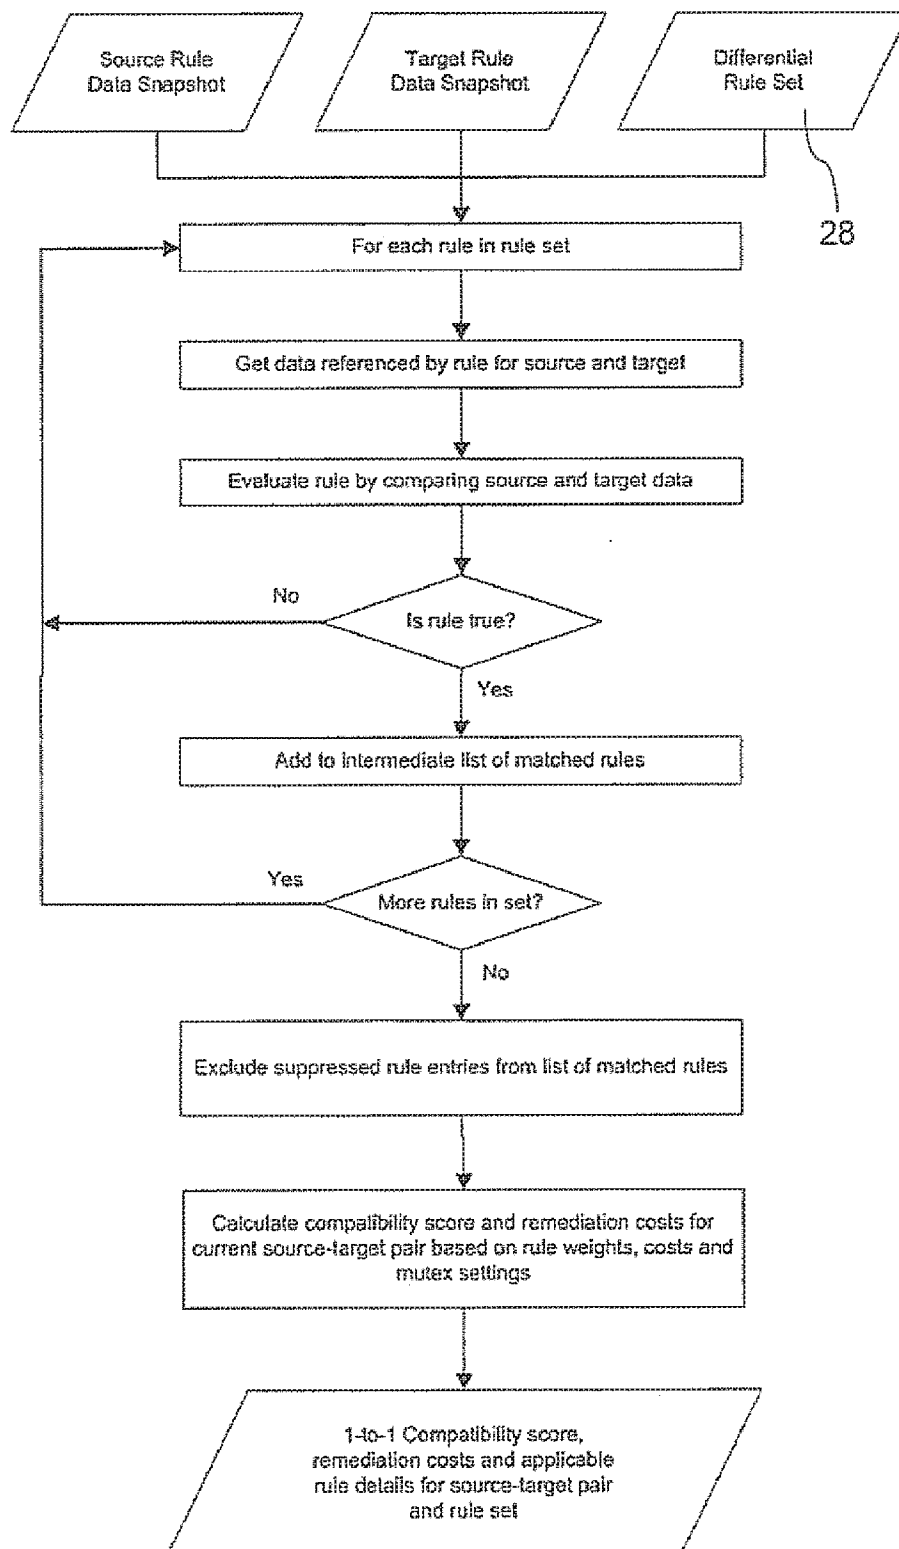


Figure 18

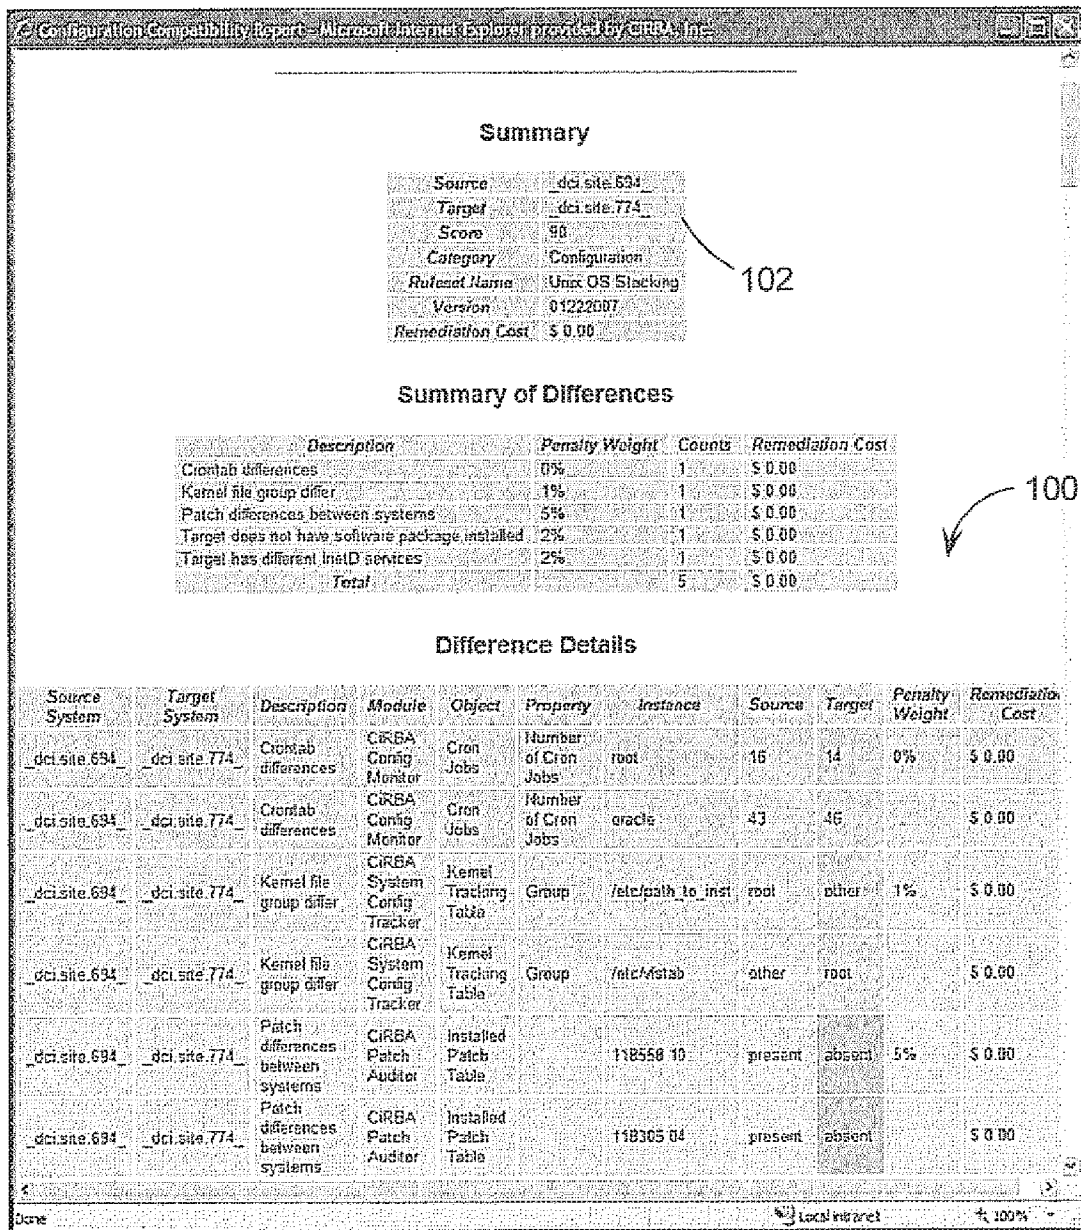


Figure 19

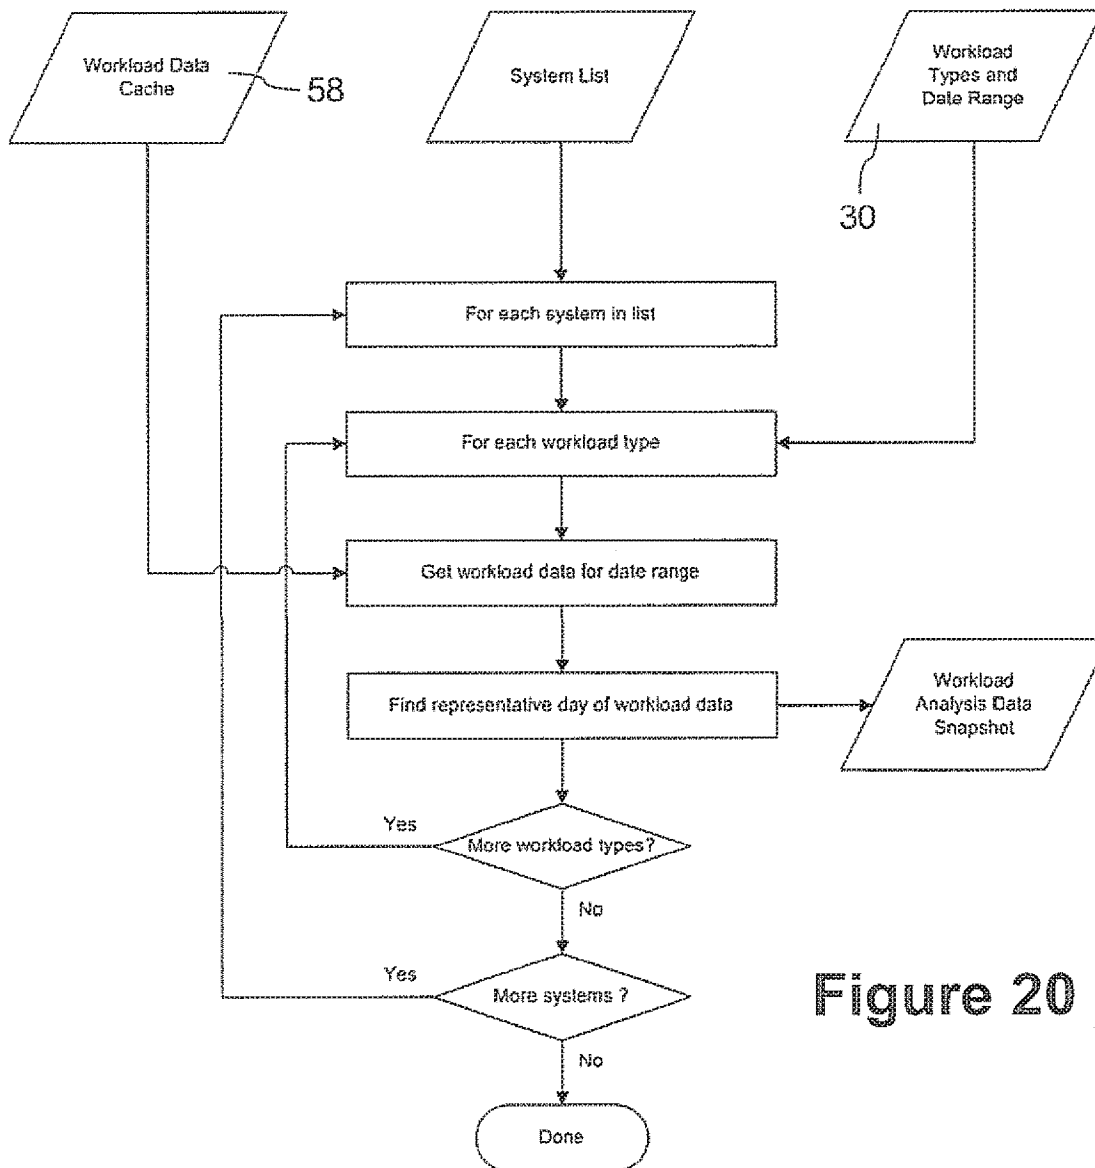
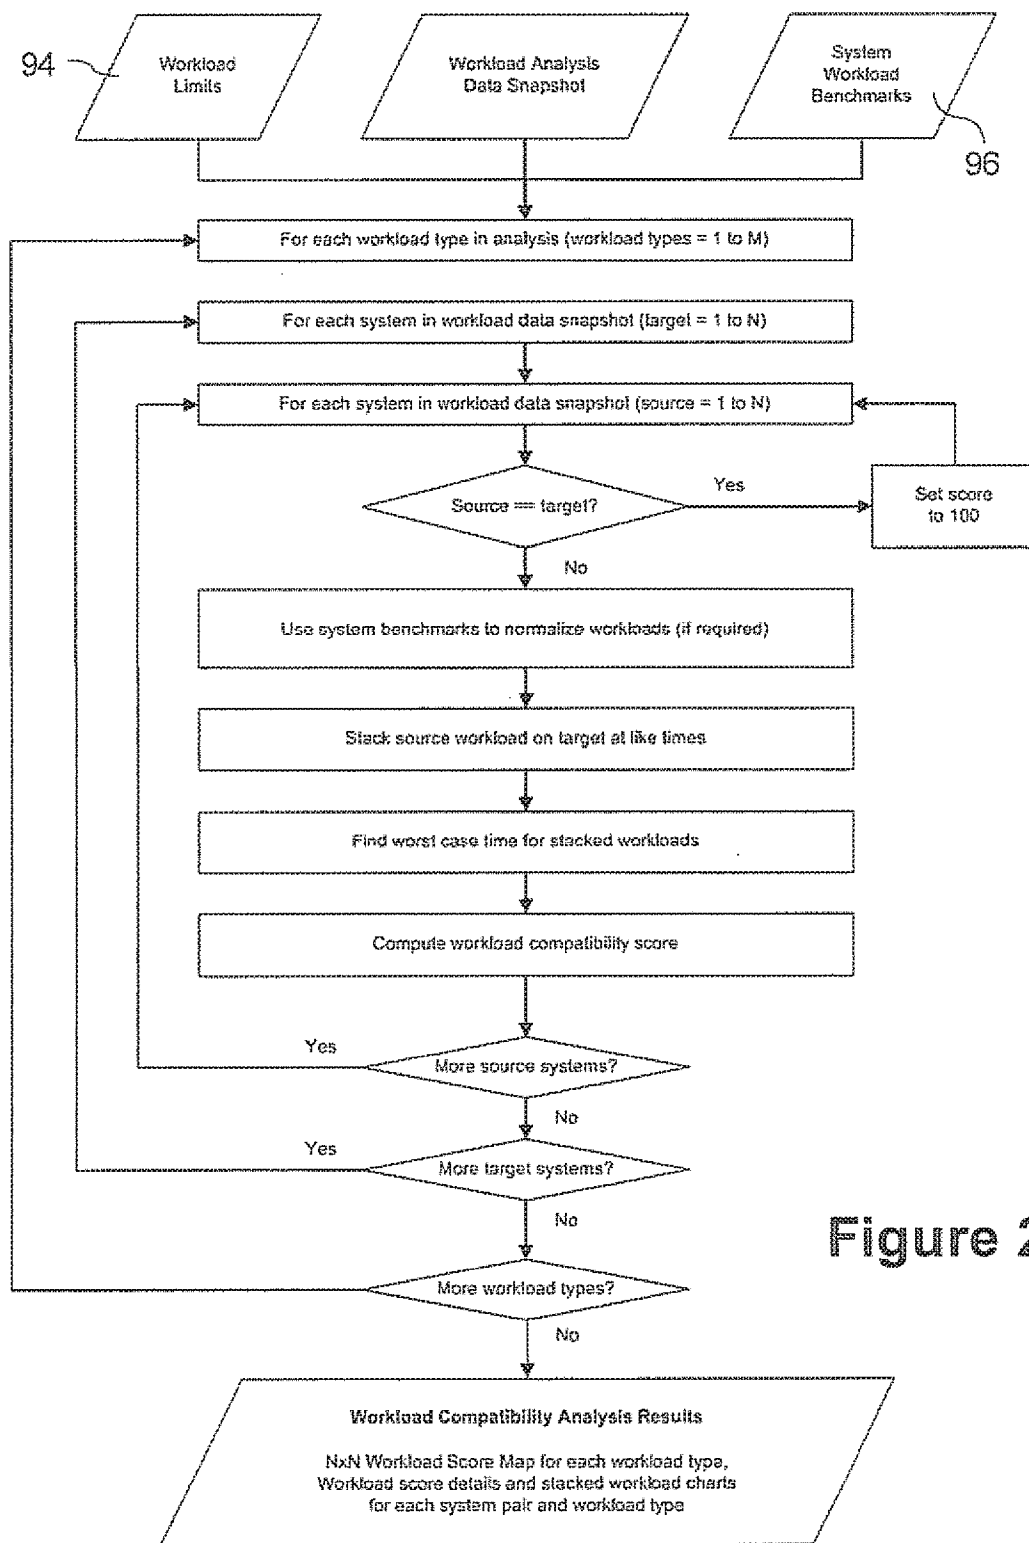


Figure 20




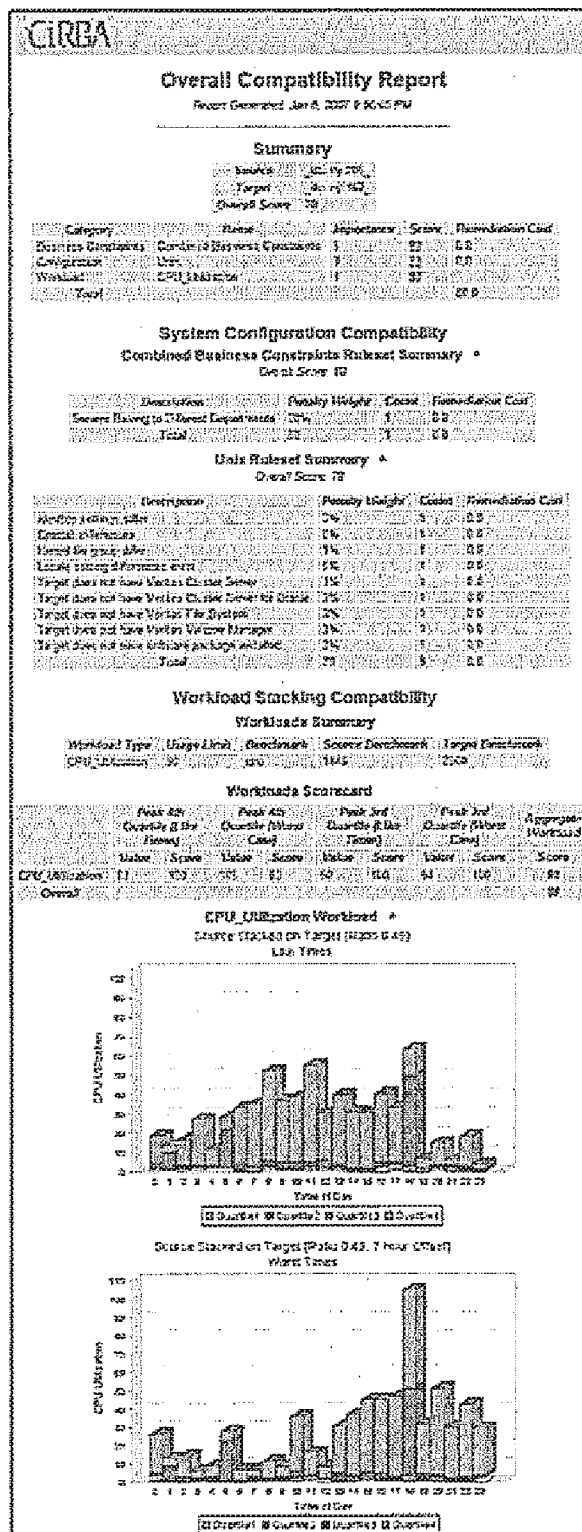


Figure 22



98

Figure 23

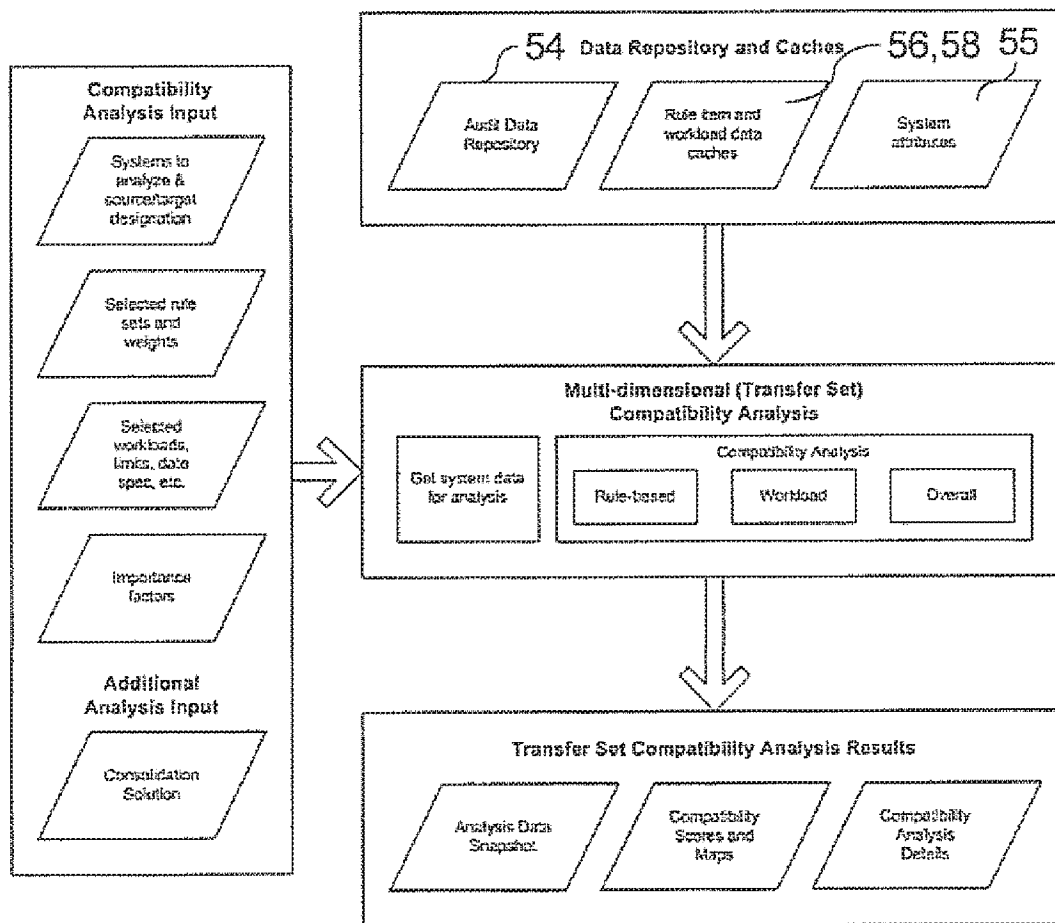


Figure 24(a)

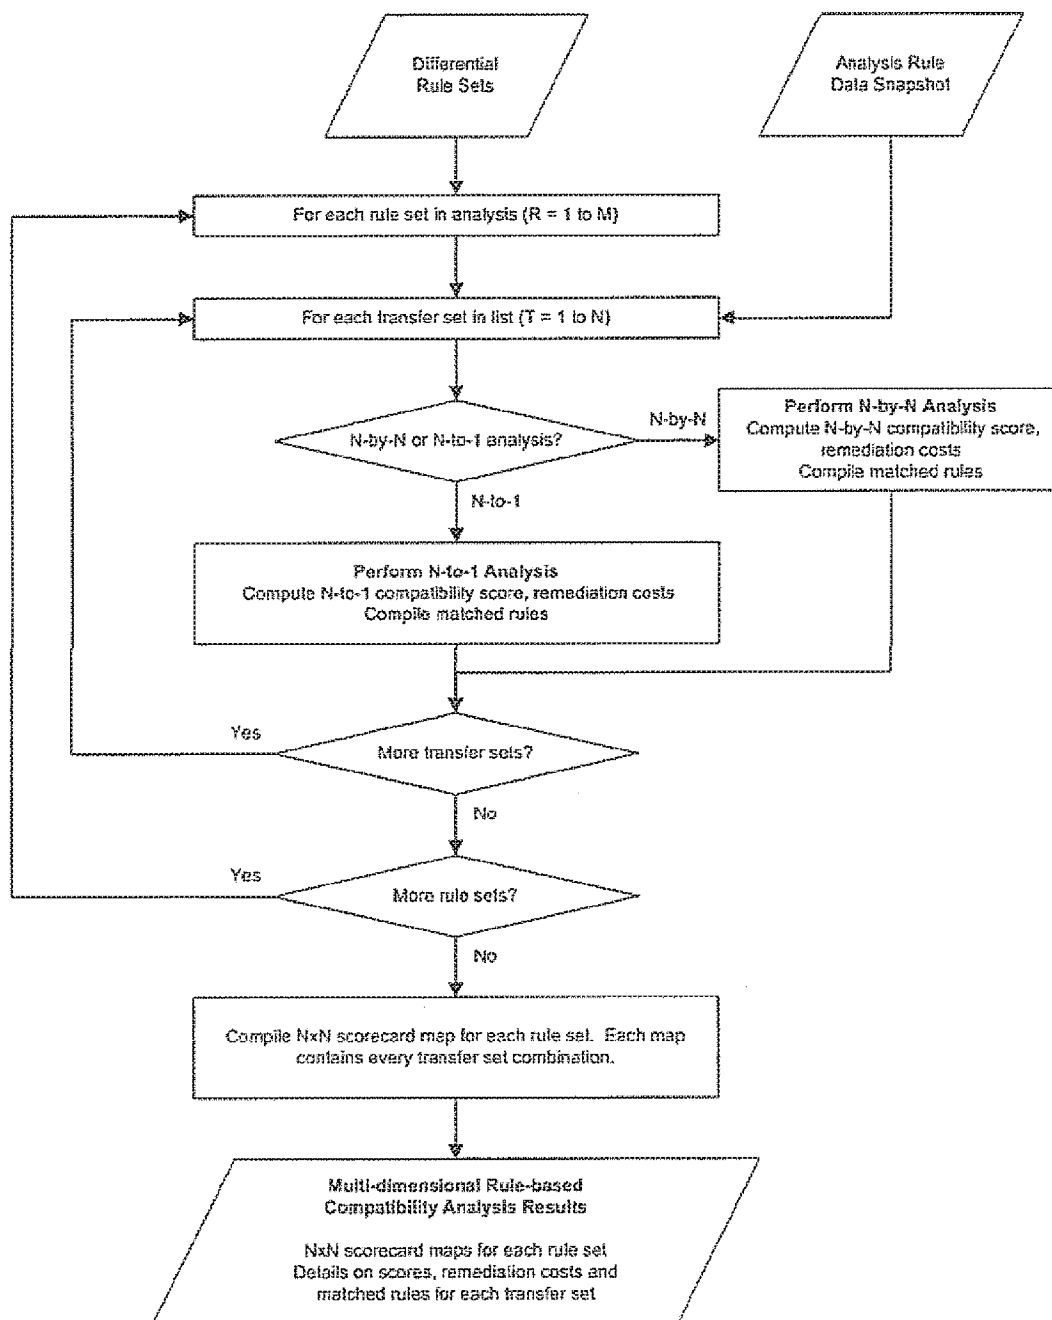


Figure 24(b)

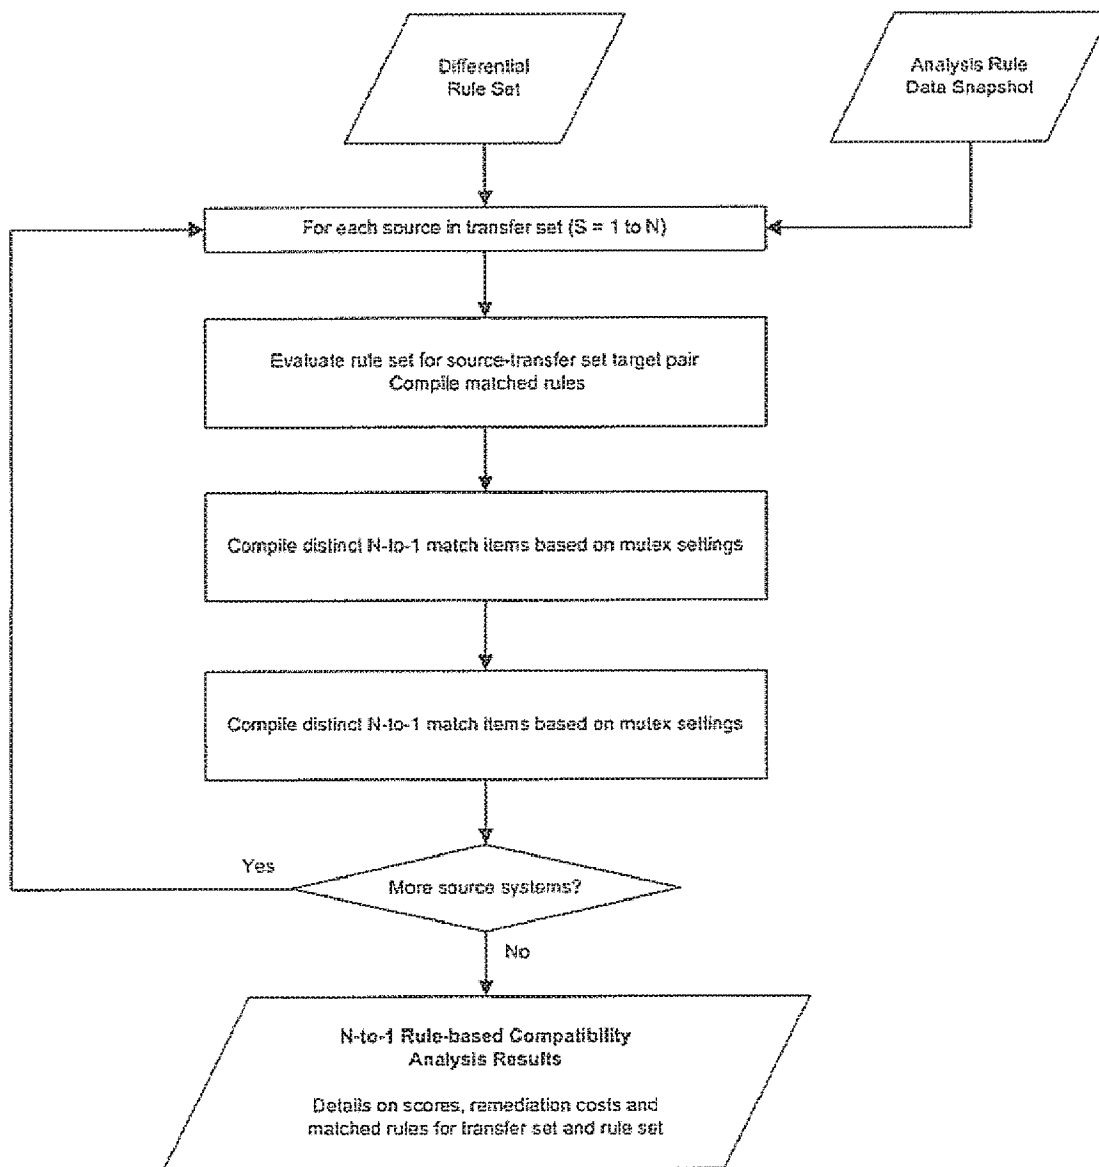


Figure 24(c)

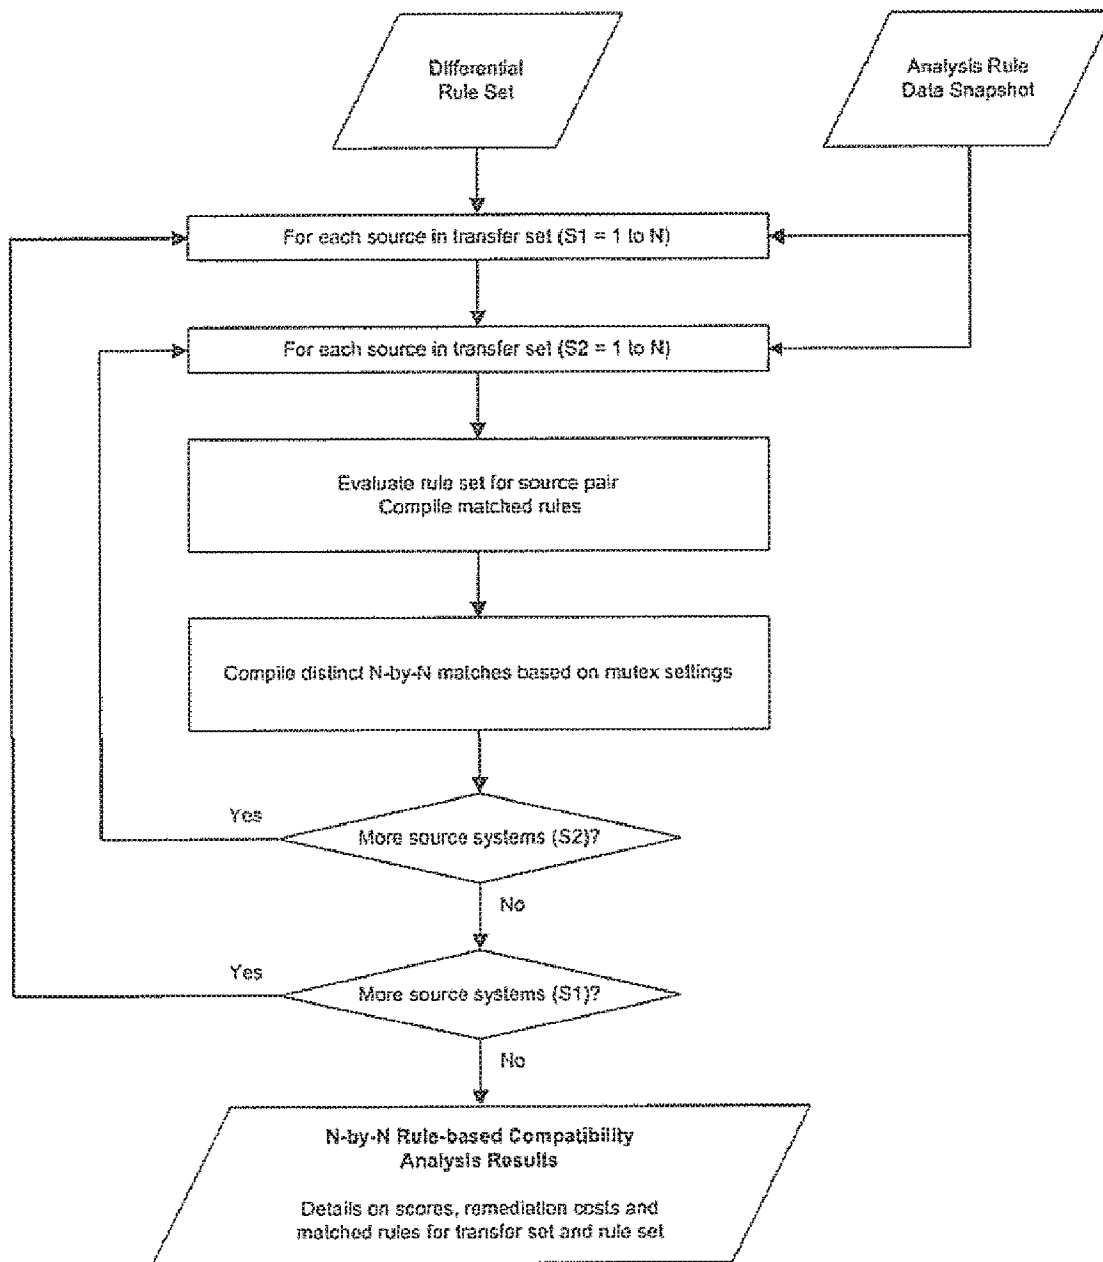


Figure 24(d)

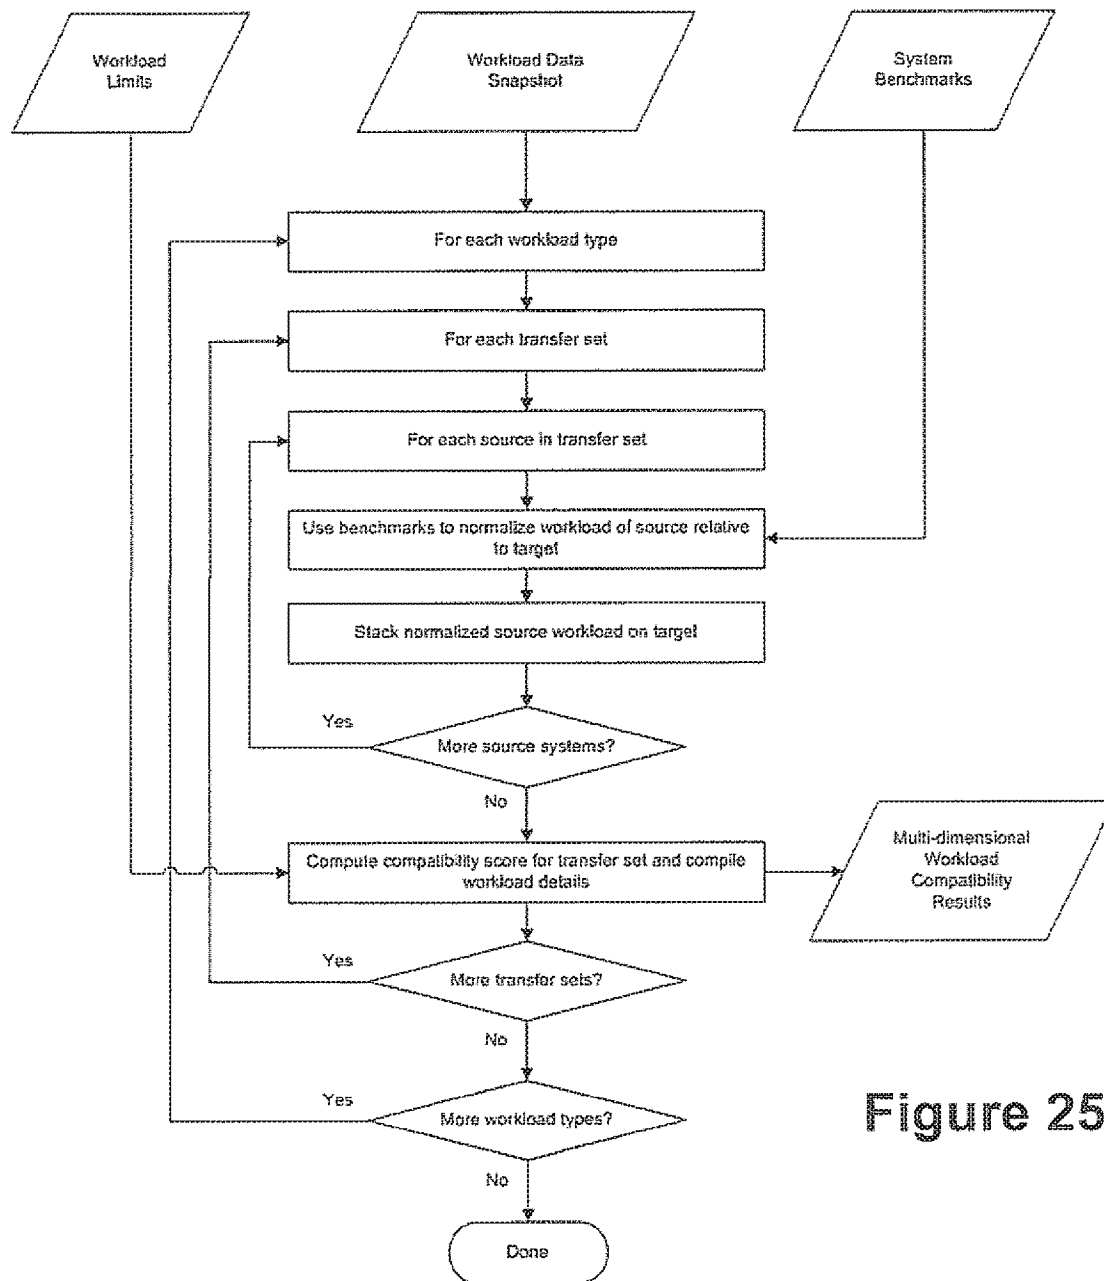


Figure 25




Figure 26

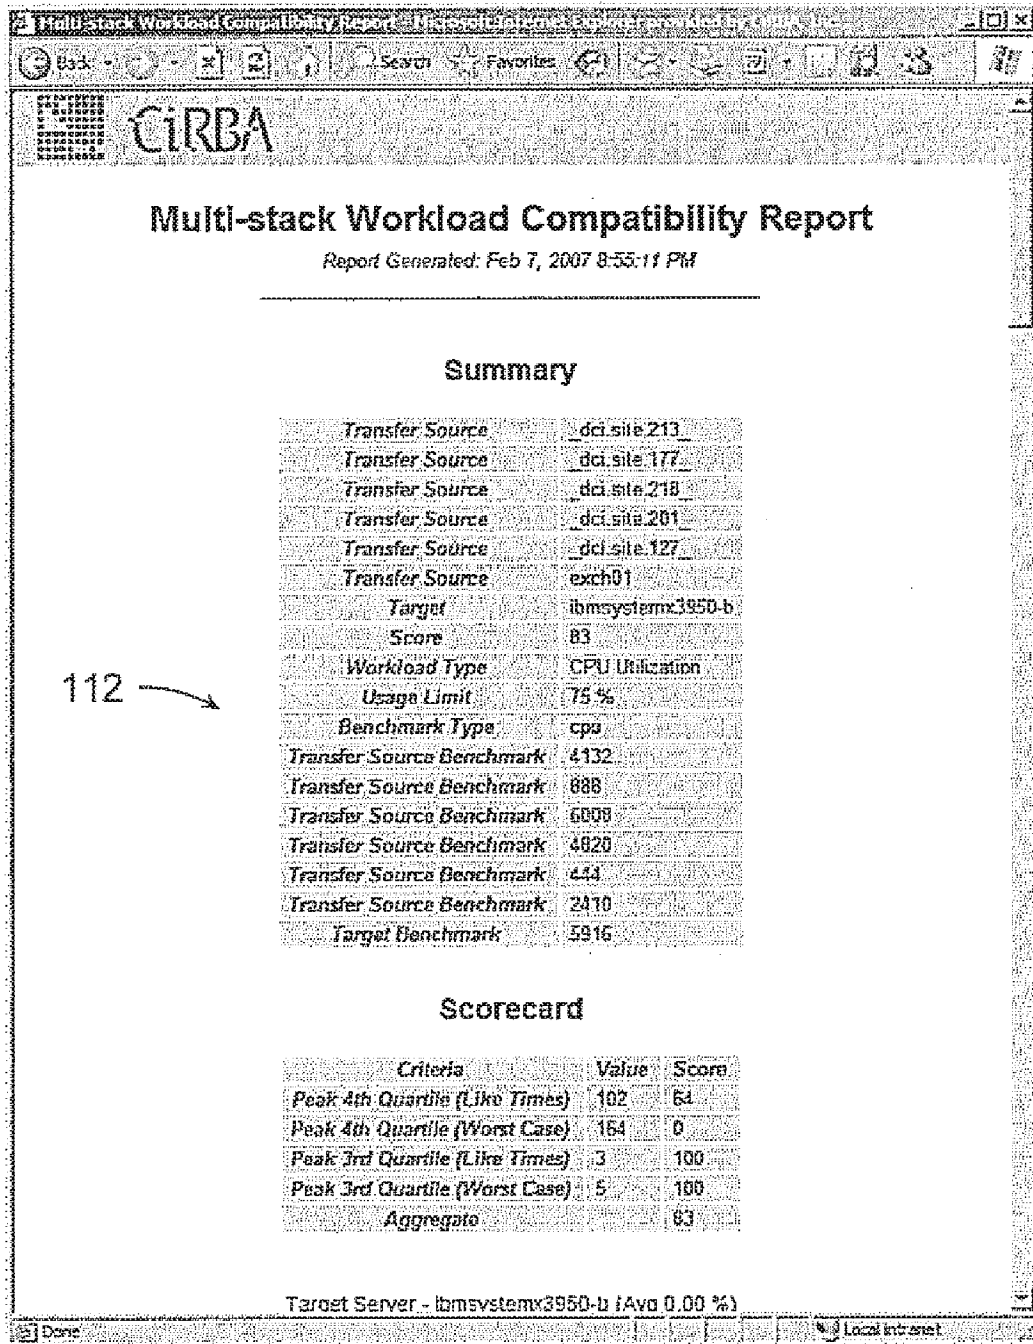


Figure 27




Figure 28

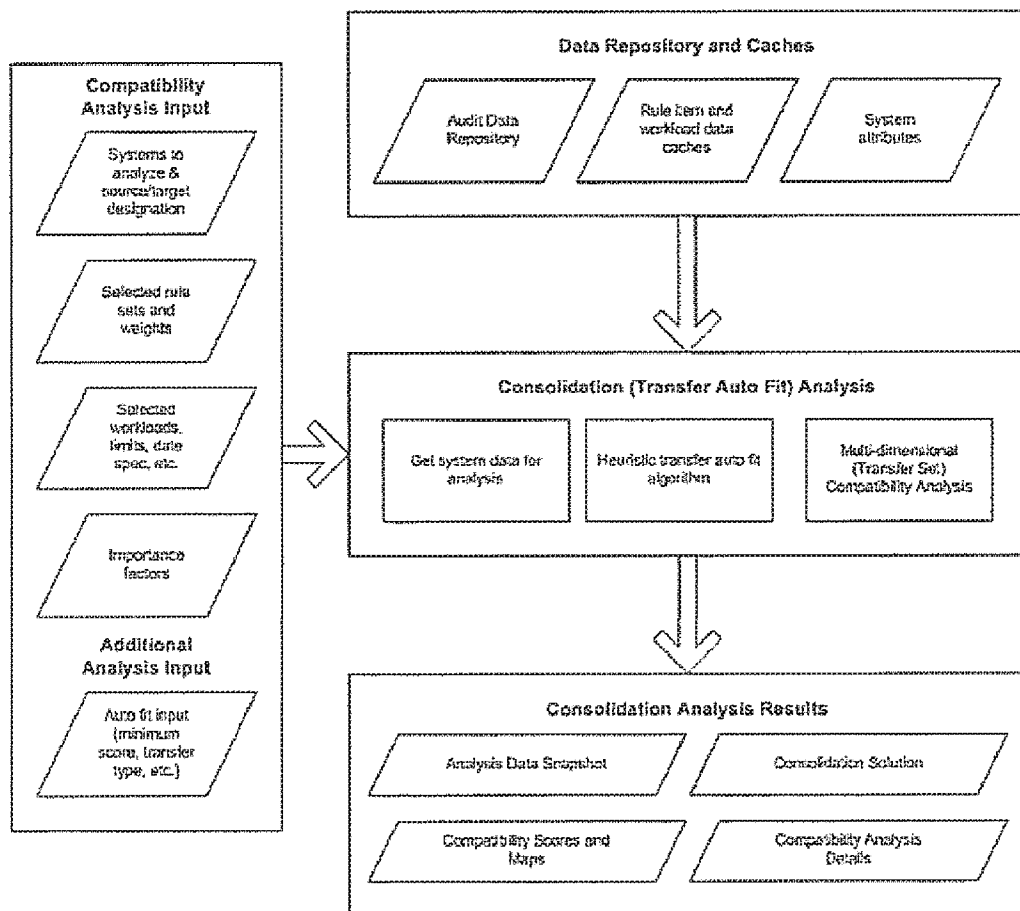


Figure 29

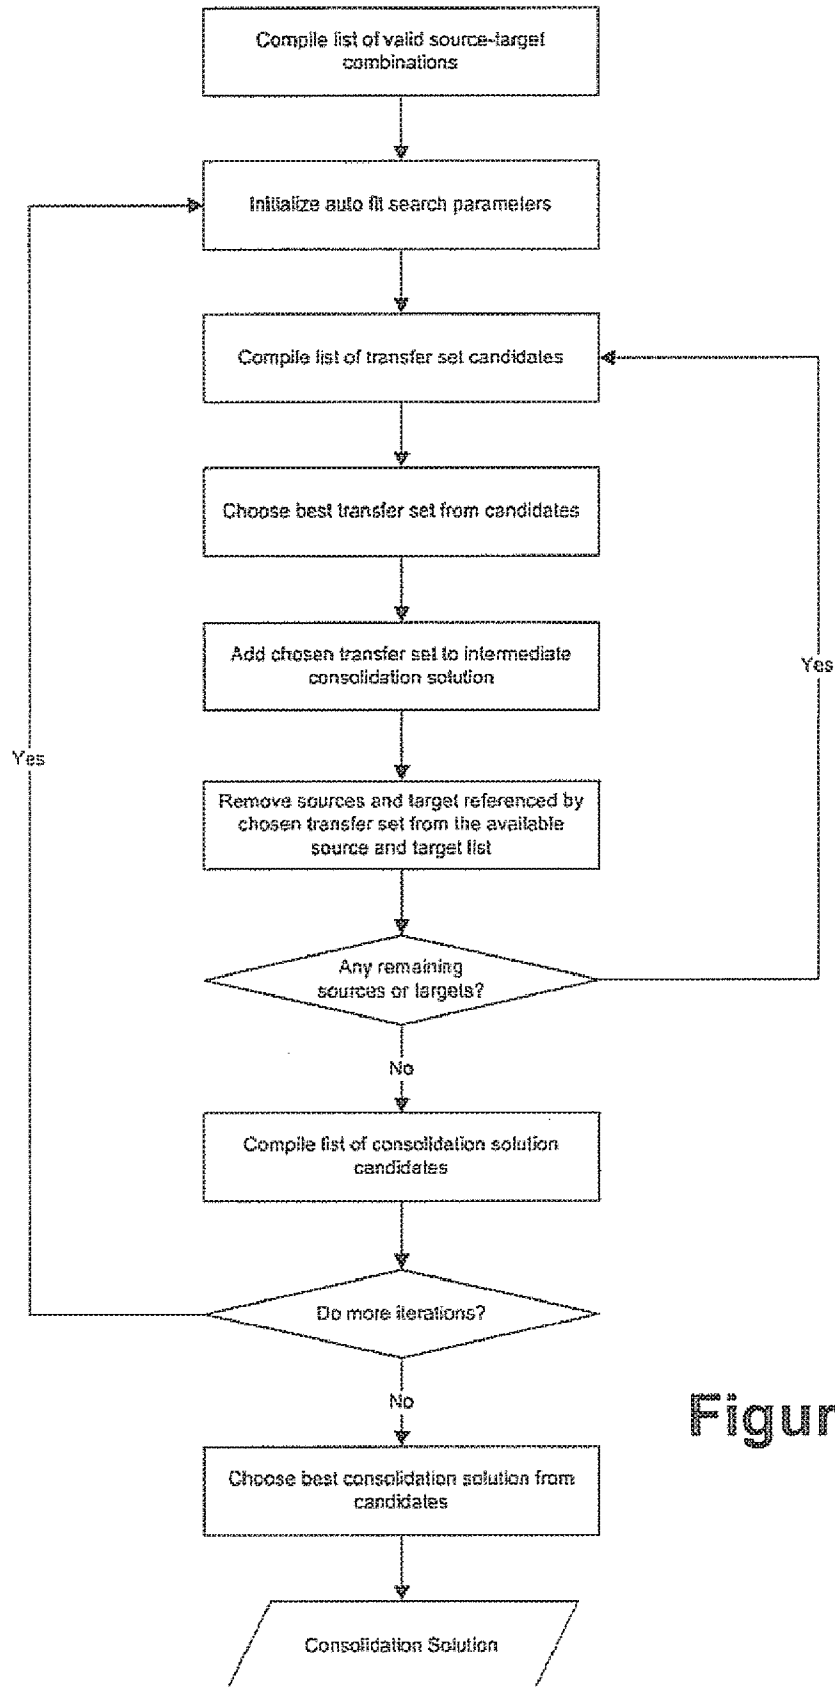


Figure 30

Analysis Summary - Microsoft Internet Explorer provided by GRBA, Inc.			
	Initial State	After Transfers	
Total Systems	27	21	
Consolidation %	-	22.22% reduction	
Consolidation % (Net of Targets)	-	22.22% reduction	
Consolidated Systems	-	6	
Transferred Systems	-	6	
Sources Only	0	0	
Targets Only	0	0	
Sources And Targets	27	21	

Rulesets			
Name	Category	Version	Importance
Unix	Configuration	1	3
Combined Business Constraints	Business Constraints	01052007	5

Workload Types		
Name	Importance	Importance
CPU_Utilization	90	1

Workload Settings	
Auto data Range	All available workload data
Workload Days Included	Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday

Transfers		
Source	Transfer	Target
dcf.ny.206	Stacks onto	_dcf.ny.110_
dcf.ny.455	Stacks onto	_dcf.ny.405_
dcf.ny.231	Stacks onto	_dcf.ny.460_
dcf.ny.329	Stacks onto	_dcf.ny.415_
dcf.ny.470	Stacks onto	_dcf.ny.435_
dcf.ny.395	Stacks onto	_dcf.ny.269_

116 ↗

Figure 31

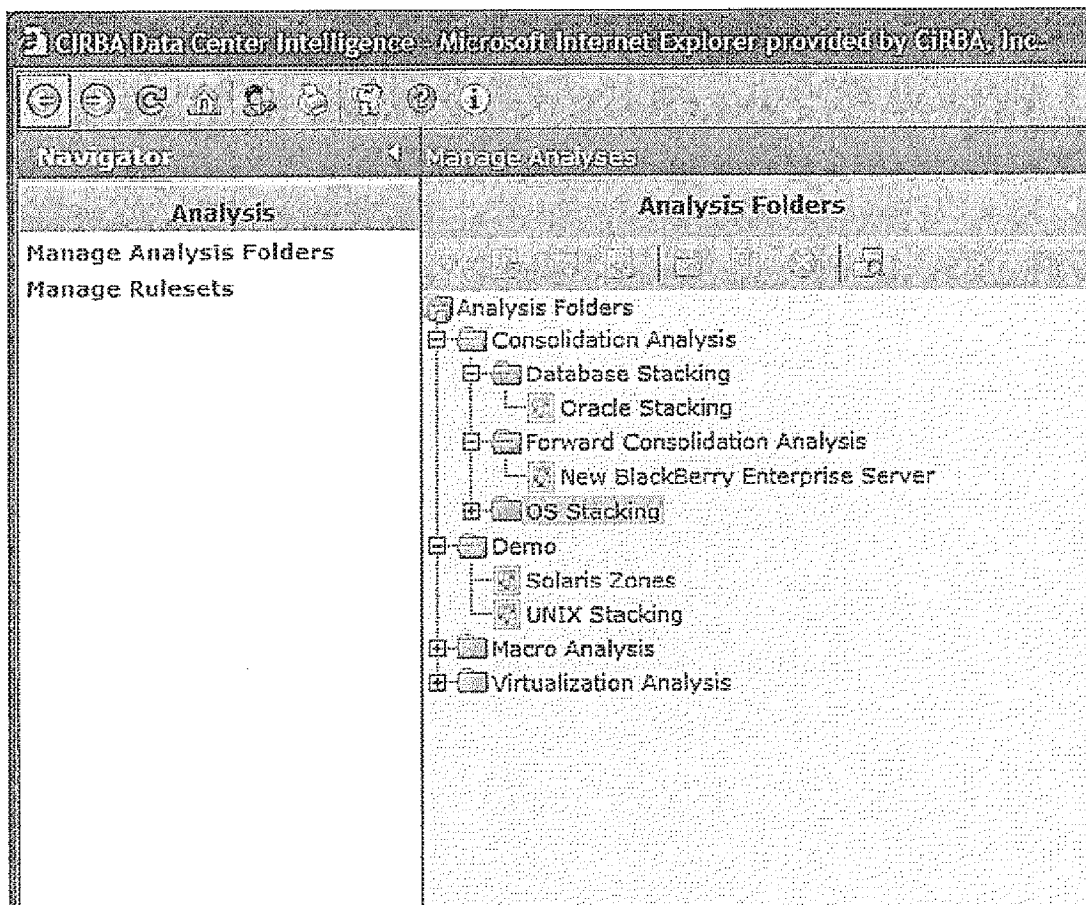


Figure 32

Demo Stacking — Web Page Dialog

Options: Importance Factors

Name: Demo Stacking

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: Configuration

Ruleset Name: Unix OS Stacking

Description: Detailed UNIX Analysis rules

Version: 0-1002007

Ruleset is User Defined

Ruleset Items

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 type="checkbox"/> Mount Mount Permissions	Target does not have CIFS mount permissions	0	0.00

Save Apply Cancel

Figure 34

Workload Settings - Web Page Dialog

Audit Date Range

☒ All available workload data

☐ Last 14 Days (Input value of 0 means today)

☐ Between 29/12/2006 and 8/1/2007 (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	90 %

Figure 35

Demo Stacking - Web Page Dialog [?] [X]

Options **Importance Factors**

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

Feature	Importance Factor	Category	Sub-category	Value
<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

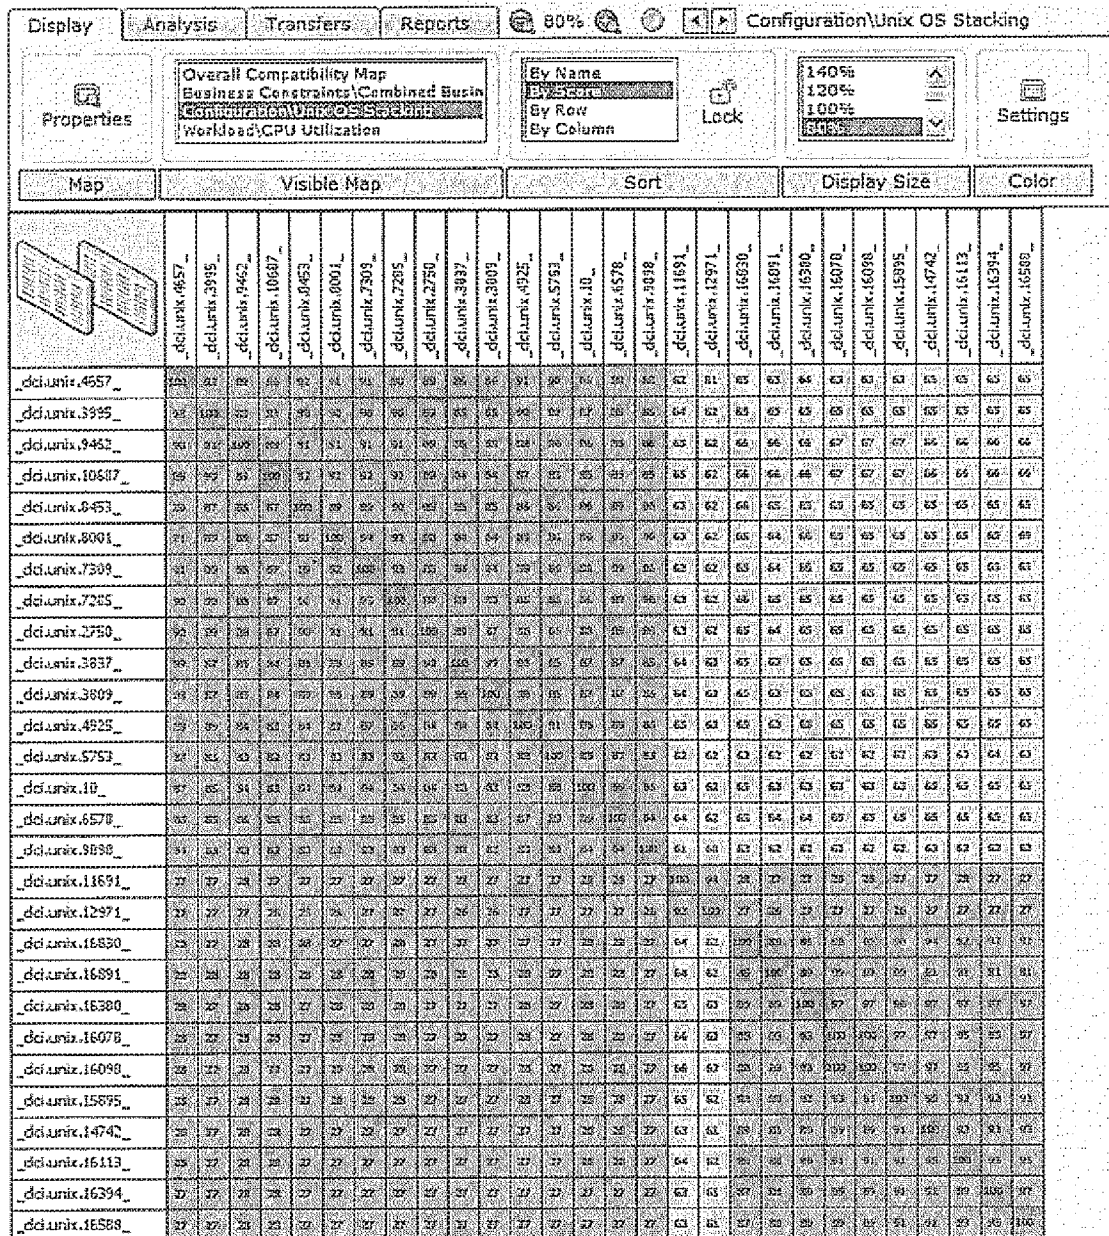


Figure 37

Configuration Compatibility Report

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

Summary

Source	pci unix 7309
Target	pci unix 11691
Score	63
Category	Configurations
Ruleset Name	Linux OS Stacking
Version	04/20/07
Remediation Cost	\$ 0.00

Summary of Differences

Description	Penalty Weight	Counts	Remediation Cost
Network settings differ	3%	1	\$ 0.00
Network settings differ	0%	1	\$ 0.00
Kernel file group differ	3%	1	\$ 0.00
Kernel file permissions differ	1%	1	\$ 0.00
Target has different kernel services	2%	1	\$ 0.00
Target has different kernel services settings	2%	1	\$ 0.00
Target running newer version of OS	30%	1	\$ 0.00
tcp_max_syn_retry settings are different	1%	1	\$ 0.00
Total		6	\$ 0.00

Difference Details

Source System	Target System	Description	Module	Object	Property	Instance	Source	Target	Penalty Weight	Remediation Cost
pci unix 7309	pci unix 11691	Network settings differ	CIRBA Solanis Details	CIRBA Solanis I/OB Table (dev/tcp)	Device Value	tcp_time_wait_interval	243000	60200	3%	\$ 0.00
pci unix 7309	pci unix 11691	Network settings differ	CIRBA Solanis Details	CIRBA Solanis I/OB Table (dev/tcp)	Device Value	tcp_max_syn_retry	2	10		\$ 0.00
pci unix 7309	pci unix 11691	Network settings differ	CIRBA Solanis Details	CIRBA Solanis I/OB Table (dev/tcp)	Device Value	tcp_max_syn_retry	162164	49112		\$ 0.00
pci unix 7309	pci unix 11691	Network settings differ	CIRBA Solanis Details	CIRBA Solanis I/OB Table (dev/tcp)	Device Value	tcp_max_syn_retry	24376	49112		\$ 0.00
pci unix 7309	pci unix 11691	Network settings differ	CIRBA Solanis Details	CIRBA Solanis I/OB Table (dev/tcp)	Device Value	tcp_synack_retries	2	1		\$ 0.00
pci unix 7309	pci unix 11691	Network settings differ	CIRBA Solanis Details	CIRBA Solanis I/OB Table (dev/tcp)	Device Value	tcp_wscale_shows	0	1		\$ 0.00

Figure 38

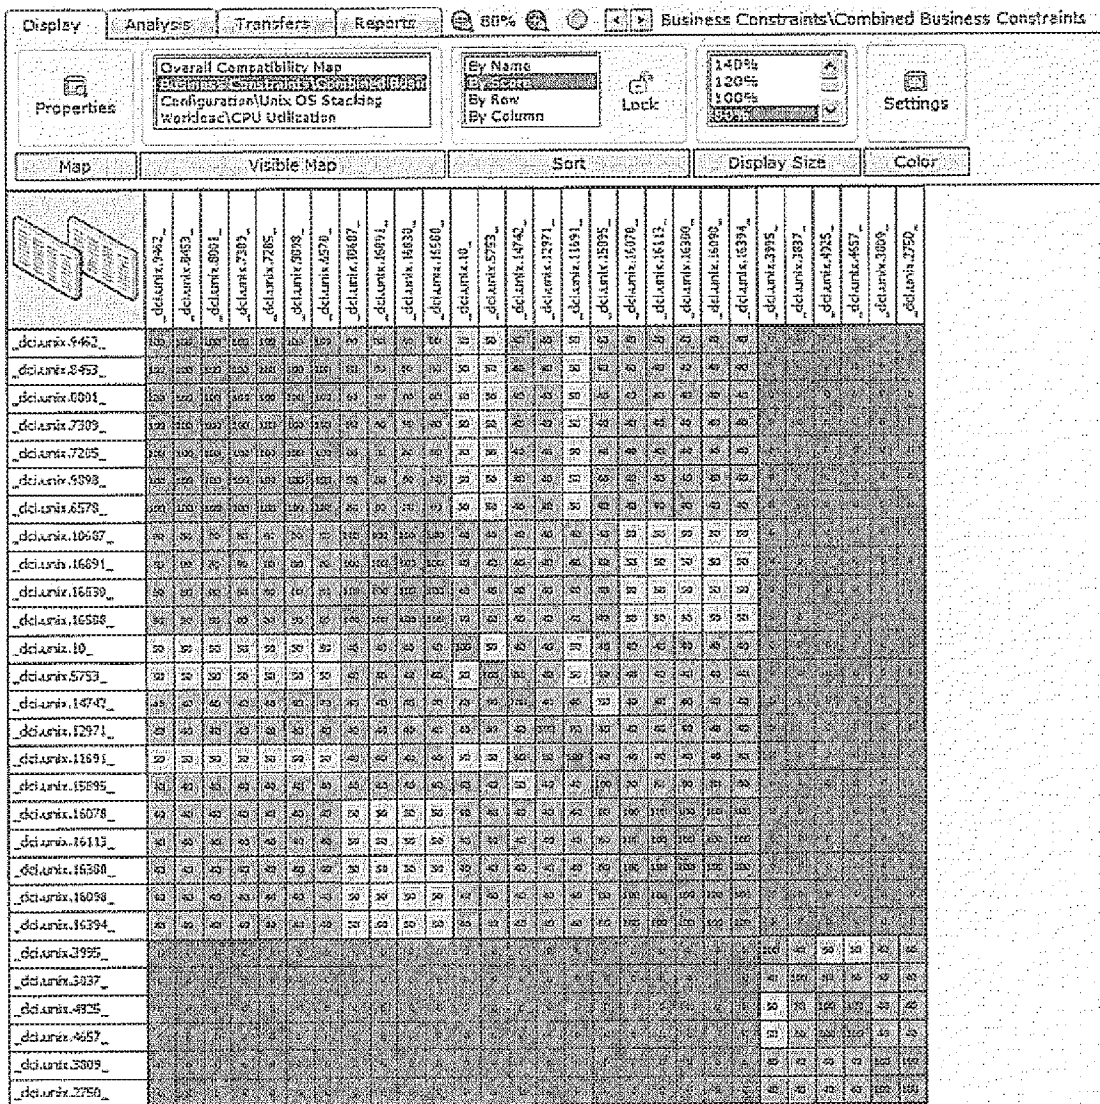


Figure 39

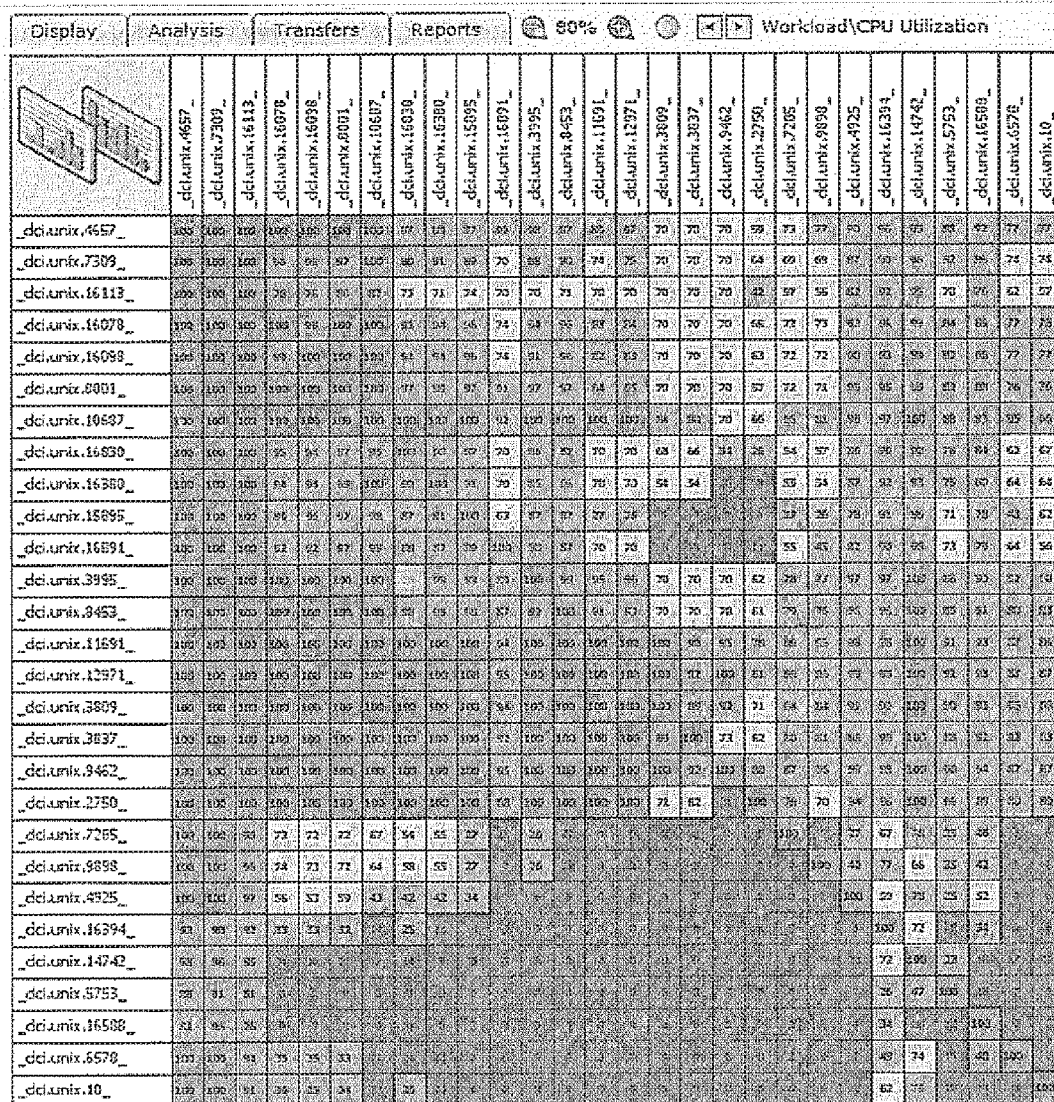


Figure 40

Workload Stacking Report

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

Summary

Source	_dc1.unix.16113_
Target	_dc1.unix.16078_
Score	78
Workload Type	CPU Utilization
Usage Limit	98 %
Benchmark Type	cpu
Source Benchmark	2872
Target Benchmark	1325

Scorecard

Criteria	Value	Score
Peak 4th Quartile (Like Times)	155	28
Peak 4th Quartile (Worst Case)	161	21
Peak 2nd Quartile (Like Times)	9	100
Peak 2nd 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	IP
Source	_dc1.unix.16113_	Solaris9.2	Jan 11, 2007 12:00:00 AM	Solaris	_dc1.unix.16049_
Target	_dc1.unix.16076_	Solaris9.2	Jan 11, 2007 12:00:00 AM	Solaris	_dc1.unix.16040_

Figure 42

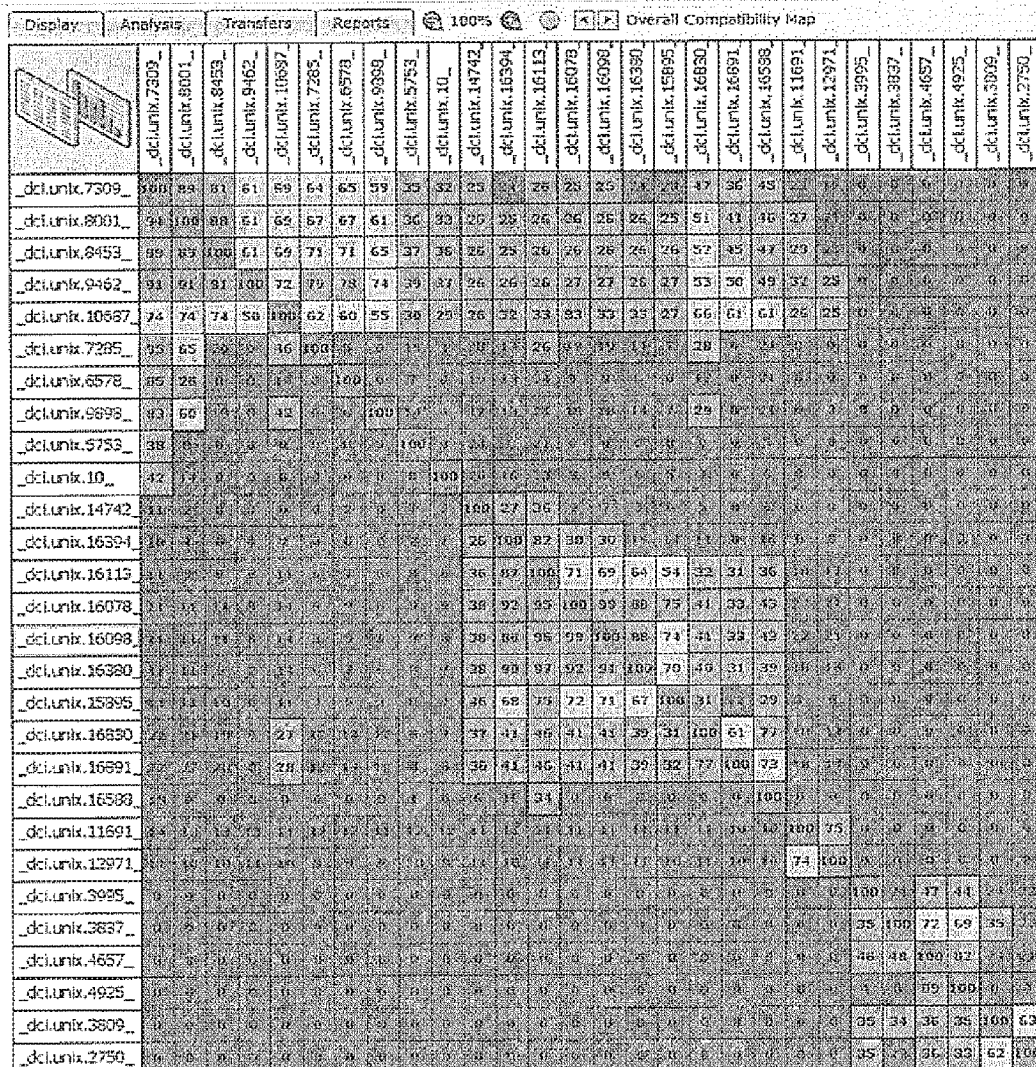


Figure 43

Overall Compatibility Report

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

Summary

Source	dcl.unix.8453
Target	dcl.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	50 %	cpi	548	101

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	Value	Score
CPU Utilization	374	0	174	0	29	100	30	100	78	78
Overall										78

CPU_Utilization Workload ^




Figure 45

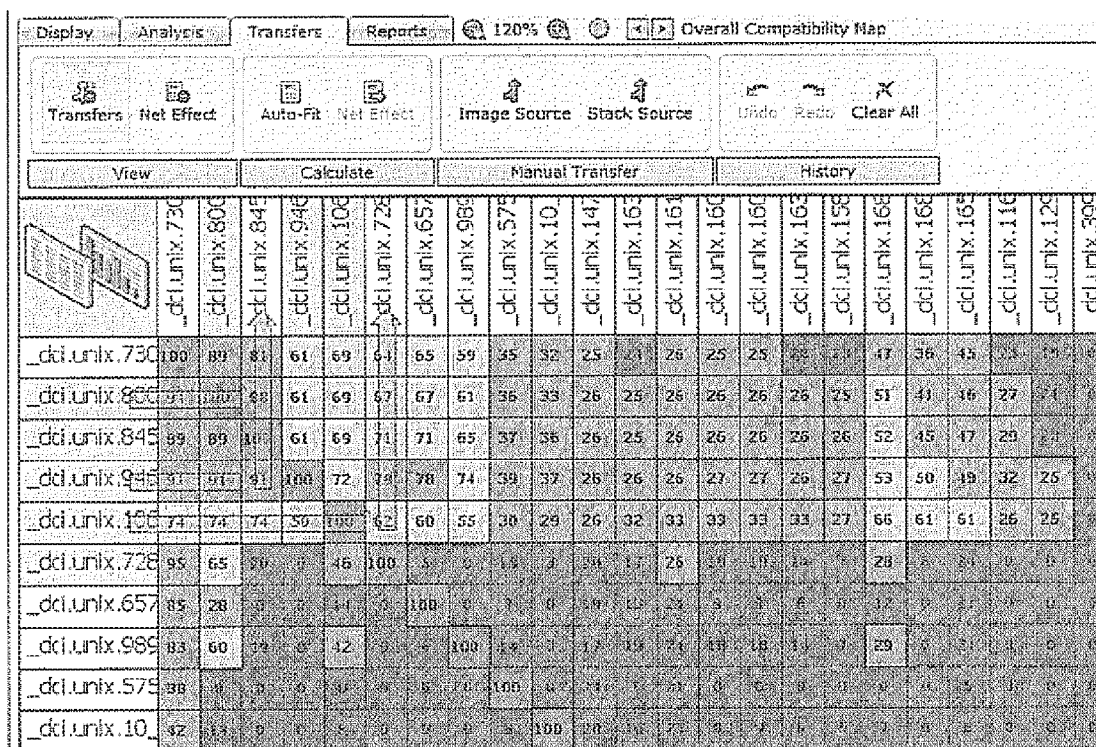


Figure 46

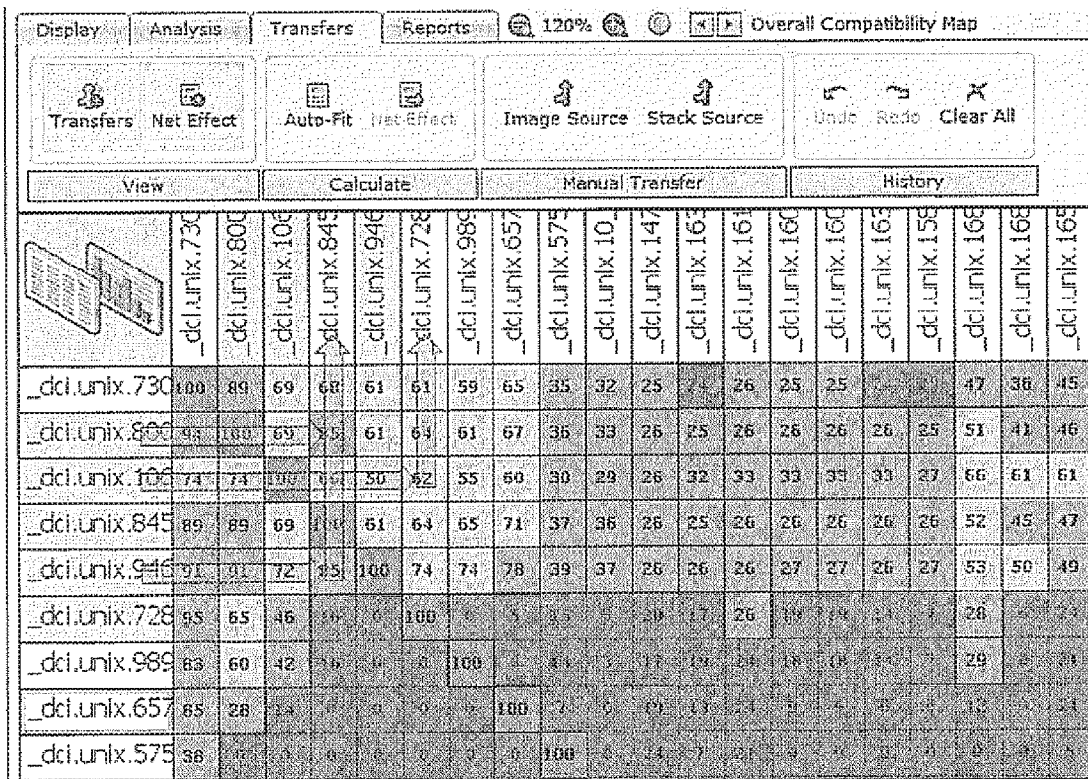


Figure 47

Overall Compatibility Report

Report Generated: Apr 21, 2007 11:51:26 AM

Summary

Transfer Source	_dcf_unix_8462_
Transfer Source	_dcf_unix_8001_
Target	_dcf_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	
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:

☒ Maximum Systems per Target:

Do not transfer if Target score is below:

☒ Clear existing Transfers

Figure 49

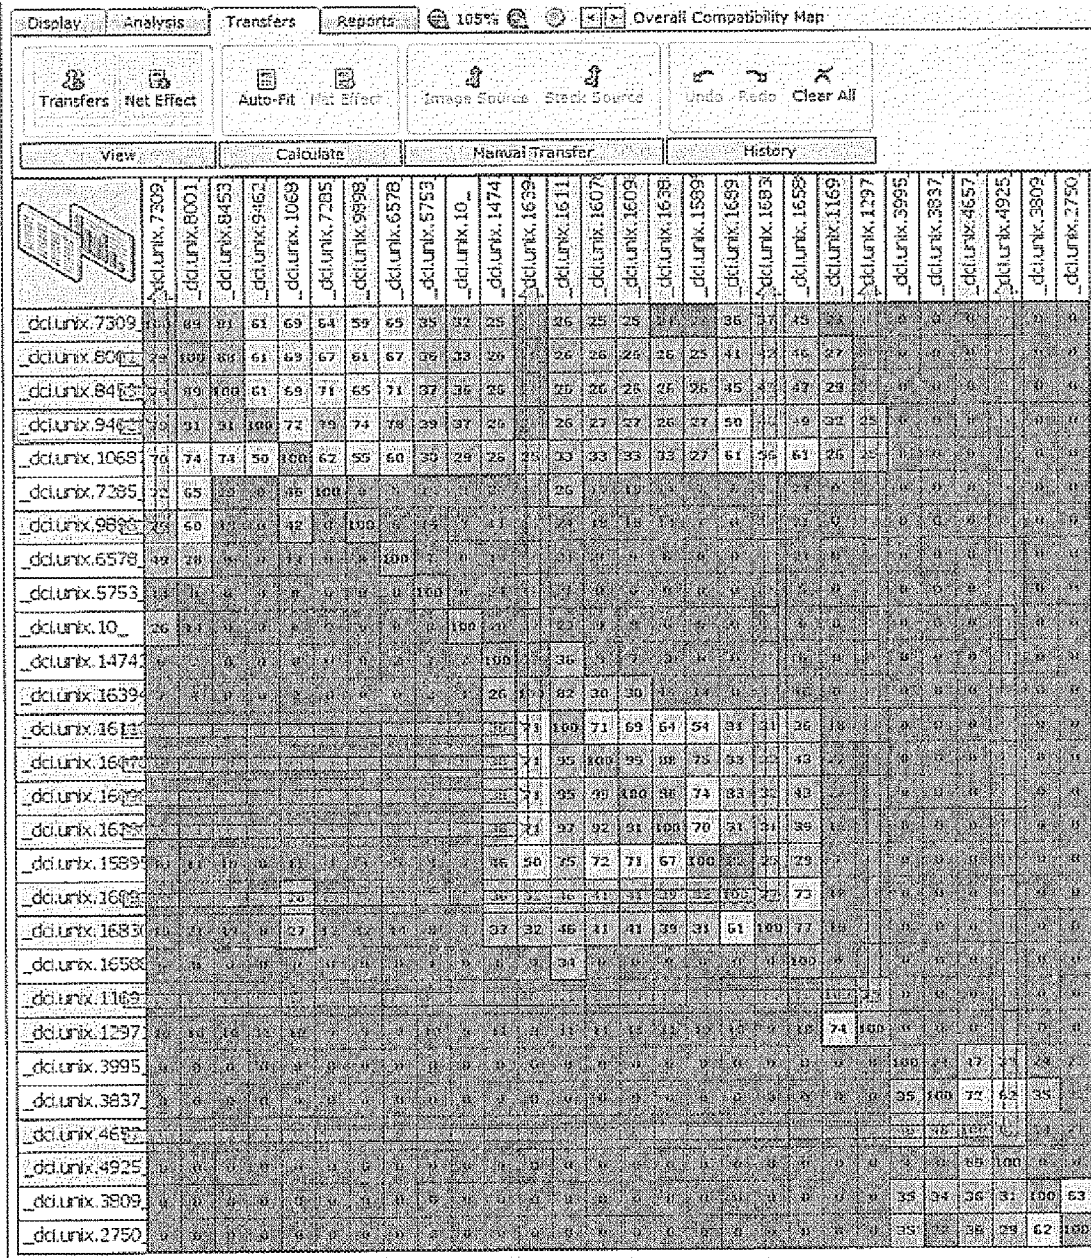


Figure 50

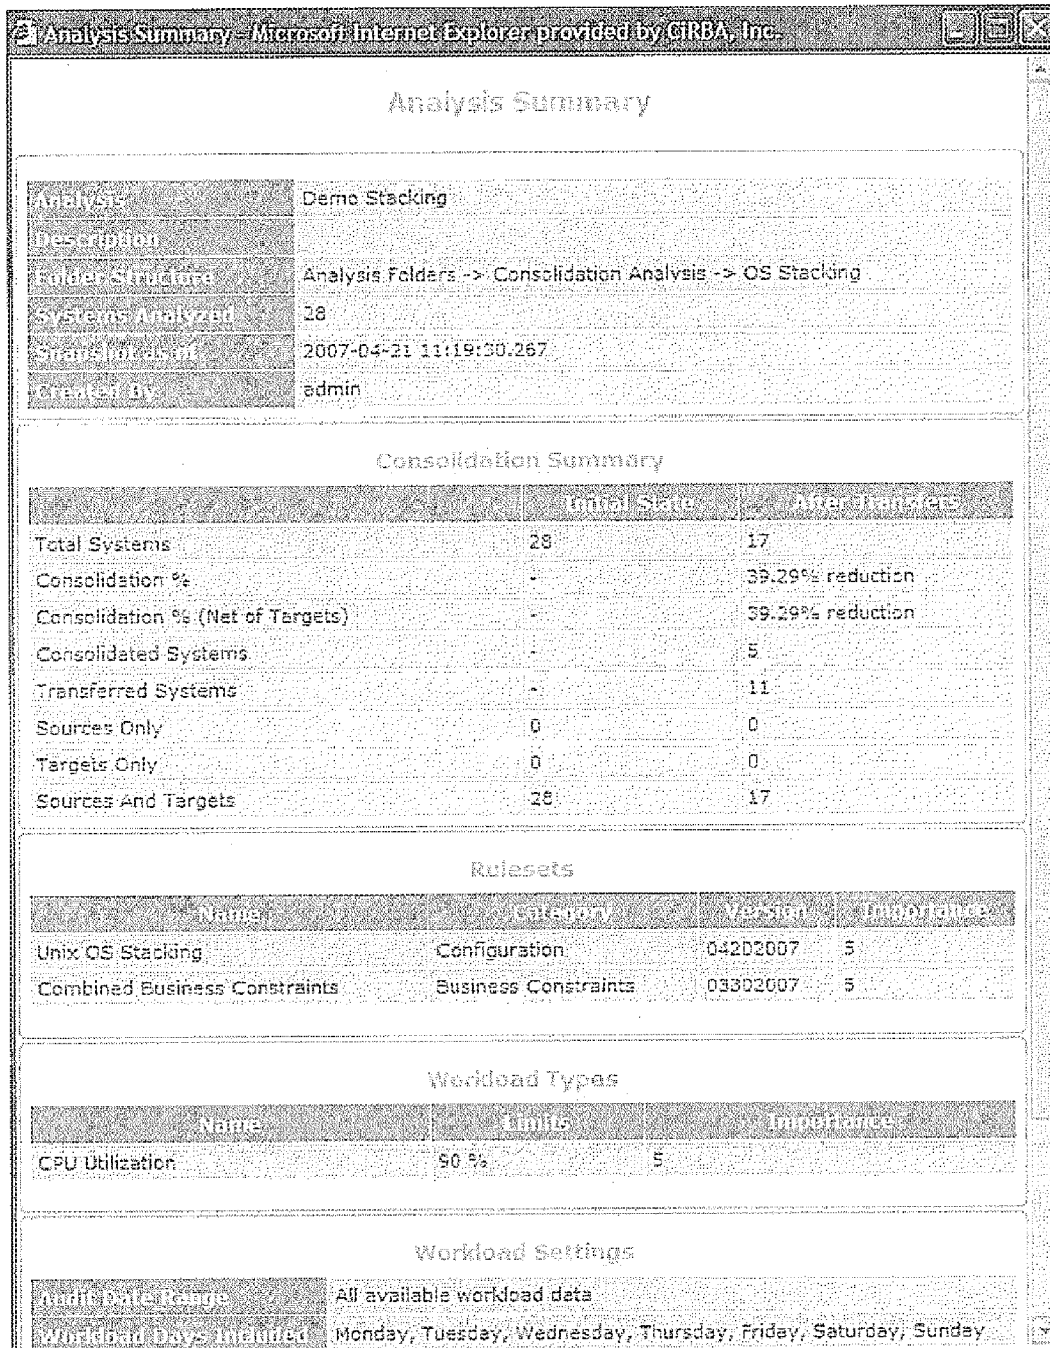


Figure 51

Transfers		
Source	Target	Target
dcf.unix.9898	Stacks onto	_dcf.unix.7309_
dcf.unix.8453	Stacks onto	_dcf.unix.7309_
dcf.unix.9462	Stacks onto	_dcf.unix.7309_
dcf.unix.8001	Stacks onto	_dcf.unix.7309_
dcf.unix.16113	Stacks onto	_dcf.unix.16394_
dcf.unix.16096	Stacks onto	_dcf.unix.16394_
dcf.unix.16380	Stacks onto	_dcf.unix.16394_
dcf.unix.16078	Stacks onto	_dcf.unix.16394_
dcf.unix.4657	Stacks onto	_dcf.unix.4935_
dcf.unix.16891	Stacks onto	_dcf.unix.16830_
dcf.unix.11691	Stacks onto	_dcf.unix.12971_

Figure 52