

1 **METHOD FOR EVALUATING COMPUTER SYSTEMS**

2 **[0001]** This application claims priority from U.S. provisional patent application
3 no. 60/745,322 filed April 21, 2006.

4 **FIELD OF THE INVENTION:**
5

6 **[0002]** The present invention relates to information technology infrastructures and has
7 particular utility in evaluating computer systems in such infrastructures.

8 **DESCRIPTION OF THE PRIOR ART**

9 **[0003]** Devices that utilize computing power such as servers, personal computers, laptops,
10 personal digital assistants (PDA) etc., are typically used in environments that require at least
11 some form of standard operation, control, communication protocol, interface etc. Such devices
12 often require upgrades, patches and security features that can change on a periodic basis.

13 **[0004]** For computing devices to communicate with each other and the supporting
14 infrastructure, they should be compatible and up to date. As organizations become more reliant
15 on computing devices of all types to perform day-to-day activities, so does the need increase to
16 periodically update and repair devices to minimize downtime and inefficiencies. Such a need
17 extends beyond central and/or distributed computing environments to mobile devices, virtual
18 networks etc.

19 **[0005]** As organizations grow and the necessary IT infrastructures also grows, the ability to
20 repair, upgrade, consolidate and evaluate computer systems becomes difficult to manage.

21 **[0006]** It is therefore an object of the following to obviate or mitigate the above-described
22 disadvantages.

23 **SUMMARY OF THE INVENTION**

24 **[0007]** In one aspect, a method of evaluating differences between a first data set and a
25 second data set for one or more computer system comprising obtaining the first data set and the
26 second data set; selecting a parameter according to a differential rule definition; comparing the
27 parameter in the first data set to the parameter in the second data set; determining if a difference

21562811.1

1 in the parameter exists between the data sets; if the difference exists, applying a weight
2 indicative of the relative importance of the difference in the parameter according to the
3 differential rule definition; and providing an evaluation of the difference according to the weight.

4 **[0008]** In another aspect, a computer readable differential rule definition for evaluating
5 differences between a first data set and a second data set for one or more computer system is
6 provided comprising a parameter for the one or more computer system; and a weight for the
7 parameter indicative of the importance of a difference in the parameter; wherein the differential
8 rule definition is used by a computer application to perform an evaluation of the difference
9 according to the weight.

10 BRIEF DESCRIPTION OF THE DRAWINGS

11 **[0009]** An embodiment of the invention will now be described by way of example only with
12 reference to the appended drawings wherein:

13 **[0010]** Figure 1 is a schematic representation of a system for analyzing computer systems.

14 **[0011]** Figure 2 is a hierarchical block diagram illustrating meta data, rules and rule sets.

15 **[0012]** Figure 3 is a schematic flow diagram showing the application of a rule set in
16 analyzing a pair of computer systems.

17 **[0013]** Figure 4 illustrates a general rule definition.

18 **[0014]** Figure 5 shows an example rule set.

19 **[0015]** Figure 6 is schematic representation of a network of systems analyzed by a computer
20 analysis program.

21 **[0016]** Figure 7 is a schematic block diagram of an underlying architecture for
22 implementing the analysis program of Figure 6.

23 **[0017]** Figure 8 is a table illustrating data enablement for system consolidation and
24 virtualization.

21562811.1

- 1 **[0018]** Figure 9 is a server compatibility index (SCI) matrix.
- 2 **[0019]** Figure 10 is an audit request template.
- 3 **[0020]** Figure 11 is a detailed configuration report and detailed configuration and workload
4 data obtained from the detailed configuration report.
- 5 **[0021]** Figure 12 is a table containing a rule set used in generating an SCI matrix.
- 6 **[0022]** Figure 13 is a screenshot of a program for generating compatibility reports.
- 7 **[0023]** Figure 14 is an SCI matrix for an example environment having four server systems.
- 8 **[0024]** Figure 15 is a table containing a summary of differences between a pair of systems
9 in the environment.
- 10 **[0025]** Figure 16 is a table containing details of the differences listed in Figure 16.
- 11 **[0026]** Figure 17 is a flowchart illustrating a system compatibility analysis procedure
12 including the application of a rule set.
- 13 **[0027]** Figure 18 is a flowchart illustrating a configuration data extraction procedure.
- 14 **[0028]** Figure 19 is a flowchart illustrating a configuration compatibility analysis procedure.
- 15 **[0029]** Figure 20 is a flowchart illustrating a rule set application procedure.

16 **DETAILED DESCRIPTION OF THE INVENTION**

- 17 **[0030]** Referring therefore to Figure 1, an analysis program 10 is in communication with a
18 set of computer systems 28 (3 are shown in Figure 1 as an example). The analysis program 10 is
19 used to evaluate differences between the computer systems 28 and provide a report showing how
20 the systems differ. The computer systems 28 may be physical systems as well as virtual systems
21 or models.

1 [0031] For the following description, a general evaluation of differences between systems
2 uses the following nomenclature: A target system refers to a system being evaluated, and a
3 baseline system is a system to which the target system is being compared. The baseline and
4 target systems may be the same system at different instances in time (baseline = prior, target =
5 now) or may be different systems being compared to each other. As such, a single system can be
6 evaluated against itself to indicate changes with respect to a datum as well as how it compares to
7 its peers.

8 [0032] In some scenarios, alternative nomenclature can be used. For example, a baseline
9 system may instead be referred to as a source system, e.g. for consolidation analyses. In such a
10 scenario, the source system is the system from which applications are moved, and the target
11 system is the system to which such applications are moved. It will be appreciated that the terms
12 "source system" and "baseline system" are herein generally synonymous, whereby a source
13 system is a type of baseline system.

14 [0033] Figure 2 illustrates a visual representation of the relationships between data used by
15 the analysis program 10. Audited data 70 is obtained from the baseline and target computer
16 systems 28 and is used to evaluate the differences between the systems 28. A distinct data set is
17 preferably obtained for each system 28 (or instance in time for the same system 28 as required).
18 Each data set comprises one or more parameter that relates to characteristics or features of the
19 respective system 28. The parameters can be evaluated by scrutinizing program definitions,
20 properties, objects, instances and any other representation or manifestation of a component,
21 feature or characteristic of the system 28. In general, a parameter is anything related to the
22 system 28 that can be evaluated, quantified, measured, compared etc.

23 [0034] Metadata 39 describes the meaning of the audited data 70 as it pertains to the
24 analysis. Preferably, comprehensive metadata 39 is included in the analysis program 10 and
25 should be capable of being modified based on the specific application and the nature of the
26 computer systems 28.

27 [0035] Differential rules 43 are conceptually a form of metadata 39 that represent the
28 importance of differences in certain parameters for the baseline and target systems 28, the

21562811.1

dependencies between different aspects of the audited data 70, and the costs associated with the remediation of differences between the system parameters.

[0036] Differential rule sets 76 are groupings of rules that represent higher-level considerations such as business objectives or administrative concerns that are taken into account when reporting on or analysing the systems 28. In this example, four differential rules 43, A, B, C and D, are grouped into two differential rule sets 76, Rule Set 1 and Rule Set 2. It will be appreciated that there may be any number of rules 43 in any number of differential rule sets 76 and those shown in Figure 2 are for illustrative purposes only.

[0037] The differential rules 43 evaluate the differences in parameters in the audited data 70 according to rule definitions. The rule definitions include weights that are indicative of the importance of the differences in particular parameters as they relates to the operation of the systems 28. The weights are applied during an evaluation of the baseline and target systems 28 if the difference exists. The evaluation may include the computation of a score or generation of other information indicative of nature of the difference(s) between the baseline and target systems 28.

[0038] The flow of data for applying an exemplary rule set 76 is shown in Figure 3. In the example shown in Figure 3, an audit engine 46 gathers audited data at step 300 from a pair of laptop computer systems 28. At step 302, a context engine 40 filters the audited data 70 using metadata 39 to determine which parameters of the laptop computer systems 28 are applicable to the analysis. A differential engine 38 and an analysis engine 41 look at the differences in parameters for the systems 28 and apply a differential rule set at step 304 which in turn evaluates the differential rule definitions for exemplary rules A, B, C and D.

[0039] As noted above, the rules 43 evaluate the differences in the baseline and target systems 28 and apply weights that indicate the importance of such differences in the parameters that have been analysed as well as the dependencies between different aspects of the data. The rule sets 76, e.g. Rule Set 1 and Rule Set 2, determine which parameters in the audited data 70 are to be evaluated and the differential rules 43 in the differential rule sets 76 are applied to the differences between the parameters in the baseline and target systems 28 based on the presence

21562811.1

1 of a difference. The difference may simply be whether or not the parameter is different but
2 nature of the difference may also be considered and have weights that vary based on how
3 different the parameter is. As such, the differential rules 43 and corresponding weights may vary
4 accordingly. For example a version 4 operating system versus a version 3 operating system may
5 be considered less costly to remedy and thus less detrimental than a version 5 operating system
6 compared to a version 1 operating system. As can be seen, even though the operating systems
7 are different in both cases, the nature of the difference can also be considered and different
8 weights and/or remedies applied accordingly.

9 **[0040]** A report generator 36 uses the results of the application of the differential rules 43 to
10 generate a report at step 306, which is then in turn displayed on the computing station 16 for
11 subsequent analysis, use and/or storage.

12 **[0041]** A general definition for a differential rule 43 is shown in Figure 4. Each rule
13 definition comprises a number rule fields and the corresponding values. A rule definition can be
14 extended to include any number of rules 43 to form a rule set 76 as shown by way of example
15 only in Figure 5. The rule definitions are computer readable and storable so that they may be
16 accessed by the program 10 and modified if necessary, for use in evaluating the computer
17 systems 28.

18 **[0042]** The rule type specifies whether the rule 43 applies to audited data directly
19 (UrlQuery) or normalized values (AliasQuery). The rule specifier specifies the URL of the data
20 object or property that is being evaluated. The optional URL fragment (i.e. the portion after the
21 “#” symbol) specifies the specific object instance (table row) that is being evaluated, with “*”
22 denoting a wildcard that matches all instances. For AliasQuery rules, this field specifies the alias
23 name.

24 **[0043]** If specified, the source field represents the literal value that would need to match the
25 value of the object/property on the source system in order for the rule 43 to match. For objects
26 and object instances, the keywords “absent” and “present” are preferably used to match cases
27 where that object is absent or present respectively. Similar to the source field, the target field
28 allows a literal match against the value of the object/property on the target system. The target

21562811.1

1 field also supports the absent/present specifiers. For numeric properties, relational operators (>, 2 <, =, !=) can be used to cause the rule 43 to trigger if the target value has the specified 3 relationship with the source value.

4 **[0044]** The weight field specifies the relative importance of that property and combination 5 of source/target values (if specified) in regard to the overall context of the comparison. Higher 6 values indicate that the condition detected by the rule 43 has a high impact on the target 7 environment.

8 **[0045]** The mutex flag field can be used to avoid multiple penalties that would otherwise 9 skew the scores. A "Y" in the mutex flag field specifies that multiple matches of the same rule 10 43 will incur only a single penalty on the overall score (as specified in the weight field), as 11 opposed to multiple accumulating penalties (which is the default behaviour).

12 **[0046]** The match flag field enables an optional symbolic flag to be "set" when a rule 43 13 matches, and which can subsequently be used to suppress other rules 43 (through the "Suppress 14 Flags" field). This effectively allows rule dependencies to be modeled in the rule set 76. The 15 suppress flag field allows symbolic flags (as specified in the "Match Flag" field) to be used to 16 suppress the processing of rules. This allows specific checks to be skipped if certain higher-level 17 conditions exist. For example, if the operating systems are different, there is no need to check 18 the patches.

19 **[0047]** The remediation cost field is preferably optional. The remediation field represents 20 the cost of "fixing" the system(s) (i.e. eliminating the condition or discrepancy detected by the 21 rule 43). When analyzing differences between (or changes to) IT systems this is used to 22 represent hardware/software upgrade costs, administrative costs and other costs associated with 23 making the required changes to the target systems. The calculations behind this field vary based 24 on the nature of the system and the parameter that would need to be added, upgraded etc.

25 **[0048]** The description field is a lay description of the condition or discrepancy detected by 26 the rule 43. These descriptions are used to provide management-level summaries when

processing rule sets 76. The description field can provide as much or as little information as required by the application.

[0049] Figure 5 provides an example rule set 76, which includes a number of rules 43. The following refers to the number indicated in the leftmost column of Figure 5.

[0050] 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 43). 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 43, 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).

[0051] 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 there are kernel parameters defined on the source that are not defined on the target. This rule is not applied if the operating systems are different.

[0052] Rule 9 scrutinizes a specific kernel setting (SHMMAX, the setting that specifies how much shared memory a system can have) and penalizes combinations where it is set to a lower value on the target than it is on the source system. Rule 10 penalizes combinations of systems that are running different versions of Oracle™. It should be noted that the remediation cost is

21562811.1

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

8 **[0053]** Rule 13 scrutinizes a windows-specific area of the audit data to determine if the
9 source and target systems are running different service pack levels. It should be noted that this
10 rule closely mirrors rule 5, which uses a rule specifier that scrutinizes the UNIX™/Linux™ area
11 of the audit data. Rule 14 scrutinizes the lists of all hotfixes applied to the source and target
12 systems and penalizes cases where they differ. This rule closely mirrors rule 6, which scrutinizes
13 patches on UNIX™ and Linux™. Rule 15 detects differing startup commands between systems.
14 Rule 16 is a rule 43 to detect differing Paths between systems, and rule 17 detects differing
15 System Paths between systems.

16 **[0054]** Rule 18 penalizes system combinations where there are services installed on the
17 source that are not installed on the target. This rule has the mutex flag set, and will therefore
18 only penalize a system combination once, no matter how many services are missing. Rule 19
19 penalizes system combinations where there are services started on the source that are not started
20 on the target. It should be noted that both the weight and the remediation cost are lower than the
21 previous rule 43, owing to the fact that it is generally easier and less expensive to start a service
22 than install it. Finally, rule 20 penalizes combinations where the target system is missing the
23 virus scanner software.

24 **[0055]** It will be appreciated that the above described rules 43 and rule set 76 are shown for
25 illustrative purposes only and that any combination of rules 43 can be used to achieve specific
26 goals. For example, rules that are applicable to the OS can be grouped together to evaluate how
27 a system 28 compares to its peers. Similarly, rules pertaining to database, Java applications etc.
28 can also be grouped.

1 [0056] The following illustrates an exemplary application of differential rules 43 for
2 analyzing compatibilities in systems 28 to according to a consolidation strategy. It will be
3 appreciated that the following is only one example of the application of a differential rule 43 and
4 should not be limited to such an example. In the following, the baseline system is referred to a
5 source system whose applications etc. are to be consolidated onto the target system.

6 [0057] Referring to Figures 6 and 7, a compatibility analysis program, generally referred to
7 by numeral 10 for clarity is deployed to gather data from the exemplary architecture shown for a
8 computing environment 12 (shown in Figure 6) and use the data with a rule set 76 to conduct an
9 evaluation of the systems 28. The analysis program 10 analyzes the environment 12 to
10 determine whether or not compatibilities exist within the environment 12 for consolidating
11 systems such as servers, desktop computers, routers, storage devices etc. The analysis program
12 10 is preferably part of a client-server application that is accessible via a web browser client 34
13 running on, e.g. a computer station 16. The analysis program 10 operates in the environment 12
14 to collect, analyze and report on audited data for not only consolidation but other functions such
15 as inventory analysis, change and compliance analysis etc. In the following examples, the
16 systems are exemplified as servers.

17 [0058] As shown in Figure 6, the example environment 12 generally comprises a master
18 server 14 that controls the operations of a series of slave servers 28 arranged in a distributed
19 system. In the example shown, the master server 14 audits a local network 18 having a series of
20 servers 28 some having local agents and others being agentless. The master server also audits a
21 pair of remote networks 20, 22 having firewalls 24. The remote network 20 includes a proxy for
22 avoiding the need to open a port range. The remote network 22 comprises a collector 30 for
23 concentrating traffic through a single point allowing an audit to be performed through the
24 firewall 24, and also comprises a proxy 32. The proxy 32 is used to convert between
25 Windows™ protocols and UNIX™/Linux™ servers, and can also concentrate traffic. The
26 proxy 32 may be required for auditing agentless Windows™ based server if the master server 14
27 is running another operating system such as UNIX™ or Linux™.

1 **[0059]** The master server 14 is capable of connecting to the slave servers 28 for performing
2 audits of configuration settings, workload etc. and thus can communicate over several applicable
3 protocols, e.g. simple network management protocol (SNMP). As shown, a computer station 16
4 running a web browser and connected to a web server (not shown) on the master server 14, e.g.
5 over HTTP, can be used to operate the analysis program 10 in the environment 12. The analysis
6 program 10 may reside on the master server 14 or may run on a remote server (not shown). The
7 analysis program 10 can gather data as it is available or retrieve a block of data from the master
8 server 14 either via electronic means or other physical means. As such, the analysis program 10
9 can operate in the environment 12 or independently (and remote thereto) so long as it can obtain
10 audited data from the environment 12. The computer station 16 enables the analysis program 10
11 to display reports and gather user input for executing an audit or analysis.

12 **[0060]** A example block diagram of the analysis program 10 is shown in Figure 7. The flow
13 of data through the program 10 begins as an audit engine 46 pulls audit data from audited
14 environments 50. The data works its way up to the web client 34 which displays an output on a
15 web interface, e.g. on computer system 16.

16 **[0061]** The audit engine 46 communicates over one or more connections referred to
17 generally by numeral 48 with audited environments 50 which are the actual systems 28, e.g.
18 server machines, that are being analysed. The audit engine 46 typically uses data acquisition
19 (DAQ) adapters to communicate with the end points (e.g. servers 28) or software systems that
20 manage the end points (e.g. management frameworks 52 and/or agent instrumentation 54). The
21 program 10 can utilize management framework adapters 52 in the audited environments 50 for
22 communicating with ESM frameworks and agent instrumentation and for communicating with
23 other agents such as a third party or agents belonging to the program 10. The audit engine 46
24 can also communicate directly with candidate and/or target systems 28 using agentless adapters
25 (central arrow in Figure 7) to gather the necessary audit information.

26 **[0062]** An audited data repository 42 is used to store audit information and previous reports.
27 The audit engine 46, using a set of audit templates 45, controls the acquisition of data that is used
28 by the other software modules to eventually generate a set of reports to display on the interface

1 34. The context engine 40 utilizes metadata 39 stored by the program 10, which indicates the
2 nature of the data, to filter out extraneous information.

3 **[0063]** The analysis engine 41 evaluates data compared in a differential engine 38 based on a
4 set of rules 43. The analysis engine 41 performs the compatibility and, in this example, the
5 consolidation analysis to determine if the environment 12 can operate with fewer systems 28.

6 **[0064]** The report generation tool 36 utilizes the set of report templates 35 for generating
7 custom reports for a particular environment 12. The report generation tool 36 utilizes the
8 information generated by the analysis engine 41. Typically, the program 10 includes a web
9 client 34 for communicating with a web interface (e.g. on computer system 16). The web
10 interface allows a user to enter settings, initiate an audit or analysis, display reports etc.

11 **[0065]** In the following examples, a source system refers to a system from which
12 applications and/or data are to be moved, and a target server or system is a system to which such
13 applications and/or data are to be moved. For example, an underutilized environment having two
14 systems 28 can be consolidated to a target system (one of the systems) by moving applications
15 and/or data from the source system (the other of the systems) to the target system.

16 **[0066]** The rules 43 and rule sets 76 can be used to evaluate systems 28 to determine
17 compatibilities based on the differences between pairs of systems 28 and the relative importance
18 of such differences to the compatibilities. The evaluation can be used for consolidation analysis,
19 compliance measures etc. For example, as system compatibility index (SCI) for each pair in a
20 plurality of systems 28 can be obtained that represents the compatibility of the systems 28 from a
21 configuration standpoint.

22 **[0067]** A system configuration compatibility analysis of N systems 18 computes NxN system
23 compatibility scores by individually considering each system 18 as a consolidation source and as
24 a target. Preferably, the scores range from 0 to 100 with higher scores indicating greater system
25 compatibility. The analysis will thus also consider the trivial cases where systems are
26 consolidated with themselves and would be given a maximum score, e.g. 100. For display and
27 reporting purposes, the scores are preferably arranged in an NxN matrix form.

21562811.1

1 [0068] An example of an SCI matrix 60 is shown in Figure 9. The SCI matrix 60 provides
2 an organized graphical mapping of system compatibility for each source/target system pair on
3 the basis of configuration data. The SCI matrix 60 shown in Figure 9 is structured having each
4 server 28 in the environment 12 listed both down the leftmost column 64 and along the
5 uppermost row 62. Each row represents a consolidation source system, and each column
6 represents the possible consolidation target. Each cell contains the score corresponding to the
7 case where the row system is consolidated onto the column (target) system.

8 [0069] The preferred output shown in Figure 9 arranges the servers 28 in the matrix such that
9 a 100% compatibility exists along the diagonal 63 where each server is naturally 100%
10 compatible with itself. The SCI matrix 60 is preferably displayed such that each cell 66 includes
11 a numerical score and a shade of a certain colour. As noted above, the higher the score (from
12 zero (0) to one hundred (100)), the higher the compatibility. The scores are pre-classified into
13 predefined ranges that indicate the level of compatibility between two systems 18. Each range
14 maps to a corresponding colour or shade for display in the matrix 60. For example, the
15 following ranges and colour codes can be used: score = 100, 100% compatible, dark green; score
16 = 75-99, highly compatible, green; score = 50-74, somewhat compatible, yellow; score = 25-49,
17 low compatibility, orange; and score = 0-24, incompatible, red.

18 [0070] The above ranges are only one example. Preferably, the ranges can be adjusted to
19 reflect more conservative and less conservative views on the compatibility results. The ranges
20 can be adjusted using a graphical tool similar to a contrast slider used in graphics programs.
21 Adjustment of the slider would correspondingly adjust the ranges and in turn the colours. This
22 allows the results to be tailored to a specific situation.

23 [0071] It is therefore seen that the graphical output of the SCI matrix 60 provides an intuitive
24 mapping between the source/target pairs in the environment 12 to assist in visualizing where
25 compatibilities exist and do not exist. In Figure 9 it can be seen that the server pair identified
26 with an asterisk (*) and by the encircled cell indicates complete compatibility between the two
27 servers for the particular strategy being observed, e.g. based on a chosen rule set. It can also be
28 seen that the server pair identified with an X and the encircled cell at the corresponding

row/column crossing comprises a particularly poor score and thus for the strategy being observed, the servers 28 in that pair are not very compatible.

[0072] The scores are calculated based on configuration data that is acquired through a configuration audit performed by the analysis program 10. The data is acquired using tools such as the table 100 shown in Figure 8 that illustrate the various types of configuration settings that are of interest and from which sources they can be obtained. In Figure 8, a number of strategies 104 and sub-strategies 105 map to various configuration sources, collectively referred to by numeral 102. Each strategy 104 includes a set of sub-strategies 105, which in turn map to specific rule sets 43.

[0073] The table 100 lists the supported consolidation strategies and the relevant data sources that should be audited to perform the corresponding consolidation analysis. In general, collecting more basis data improves the analysis results. The table 100 enables the analysis program 10 to locate the settings and information of interest based on the strategy 104 or sub-strategy 105 (and in turn the rule set) that is to be used to evaluate the systems 28 in the environment 12. The results can be used to determine source/target candidates for analysing the environment for the purpose of, e.g. consolidation, compliance measures etc.

[0074] The score provided in each cell indicates the configuration compatibility for consolidating pairs of servers. The matrix 60 provides a visual representation of the compatibilities and an intuitive way to evaluate the likelihood that systems can be consolidated and have associated tools (as explained below) that can be used to analyse compliance and remediation measures to modify systems 28 so that they can become more compatible with other systems 28 in the environment 12. It can therefore be seen that a significant amount of quantitative data can be analysed in a convenient manner using the graphical matrix 60 and associated reports and graphs (described below).

[0075] For example, a server pair that is not compatible only for the reason that certain critical software upgrades have not been implemented, the information can be uncovered through analysis tools used with the SCI matrix 60, and then investigated, so that upgrades can be implemented, referred to herein as remediation. Remediation can be determined by modeling

21562811.1

1 cost of implementing upgrades, fixes etc that are needed in the rule sets. If remediation is then
2 implemented, a subsequent analysis may then show the same server pair to be highly compatible
3 and thus suitable candidates for consolidation.

4 **[0076]** The SCI matrix 60 can be sorted in various ways to convey different information. For
5 example, sorting algorithms such as a simple row sort, a simple column sort and a sorting by
6 group can be used.

7 **[0077]** A simple row sort involves computing the total scores for each source system (by
8 row), and subsequently sorting the rows by ascending total scores. In this arrangement, the
9 highest total scores are indicative of source systems that are the best candidates to consolidate
10 onto other systems.

11 **[0078]** A simple column sort involves computing the total scores for each target system (by
12 column) and subsequently sorting the columns by ascending total score. In this arrangement, the
13 highest total scores are indicative of the best consolidation target systems.

14 **[0079]** Sorting by group involves computing the difference between each system pair, and
15 arranging the systems to minimize the total difference between each pair of adjacent systems in
16 the matrix. The difference between a system pair can be computed by taking the square root of
17 the sum of the squares of the difference of a pair's individual compatibility score against each
18 other system in the analysis. In general, the smaller the total difference between two systems,
19 the more similar the two systems with respect to their compatibility with the other systems. The
20 group sort promotes the visualization of the logical breakdown of an environment by producing
21 clusters of compatible systems 18 around the matrix diagonal. These clusters are indicative of
22 compatible regions in the environment 12.

23 **[0080]** The following illustrates an example for generating the SCI matrix 60 discussed
24 above. Figure 17 shows generally a configuration compatibility analysis. A configuration data
25 extraction process is performed using per-system configuration data and a compatibility rule set.
26 The configuration data extraction process produced filtered per-system configuration data and

1 this filtered data is used to perform the configuration compatibility analysis to in turn produce
2 system compatibility results, e.g. including SCI scores in a SCI matrix 60.

3 **[0081]** The configuration data extraction step is shown in greater detail in Figure 18. The
4 per-system configuration data comprises a data set for each system 28 that is obtained during the
5 auditing process. The compatibility rule set defines which settings are important for determining
6 compatibility. As discussed above, the compatibility rule set is typically a predefined set of rules
7 which can be revised as necessary based on the specific environment 12. The rule set is thus
8 preferably compiled according to the target systems being analysed and prior knowledge of what
9 makes a system compatible with another system for a particular purpose.

10 **[0082]** Configuration data extraction analyses the per-system configuration data and the
11 compatibility rule sets 76. Each system is analysed for each rule 43 in the rule set 76. For each
12 rule, the configuration data referenced by the rule (according to its definition) is extracted and
13 saved by the analysis engine 41. The extraction process results in a filtered data set for each
14 system 28 that corresponds to the actual configuration data that can be used to determine
15 compatibilities between the systems 28. The filtered data is used for the compatibility analysis.

16 **[0083]** An exemplary configuration compatibility analysis procedure is shown in Figures 19-
17 20. When analyzing system compatibility, the list of target and source systems 28 are the same.
18 The compatibility is evaluated in two directions, e.g. for a Server A and a Server B, migrating A
19 to B is considered as well as migrating B to A.

20 **[0084]** Turning first to Figure 19, for each target system T (T = 1 to N where N is the number
21 of systems), the differential engine 38 first looks at each source system S (S = 1 to N). If the
22 source=target then the SCI score for that source is set to 100, no further analysis is required and
23 the next pair is analyzed. If the source and target are different, the rules are evaluated against the
24 source/target pair to compute the SCI score, remediation cost and to compile the associated rule
25 details. Estimated remediation costs are optionally specified with each rule item. As part of the
26 rule evaluation and subsequent SCI score calculation, if a rule is true (a difference is identified),
27 the corresponding cost to address the deficiency is added to the remediation cost for the pair of
28 systems 28 being analysed.

21562811.1

1 [0085] The evaluation of the rules is shown in Figure 20. The evaluation of the rules
2 considers the filtered configuration data for both the source system and the target system, as well
3 as the compatibility rule set that is being applied. For each rule in the set, the data referenced by
4 the rule is obtained from both the target data and source data. The rule is evaluated by having
5 the differential engine 38 compare the data. If the rule is not true (i.e. if the systems are the same
6 according to the rule definition) then the data is not considered in the SCI score and the next rule
7 is evaluated. If the rule is true, the rule details are added to an intermediate result. The
8 intermediate result includes all true rules.

9 [0086] Preferably, a suppression tag is included with applicable rules. As previously
10 discussed, the suppression tag indicates other rules that are not relevant if that rule is true. For
11 example, if the OS in a source/target pair is different, there is no need to check whether or not
12 the patches are the same, since different OSs should invariably have different patches. The
13 suppression flag allows the program 10 to avoid unnecessary computations. As also previously
14 discussed, a mutex flag can be used to avoid unfairly reducing the score for each true rule when
15 the rules are closely affected by each other. For example, if several patches are different, the
16 score is only docked marks once, for that rule type so as to not diminish the compatibility score
17 due to only one difference seen multiple times.

18 [0087] Once each rule has been evaluated, a list of applicable rules is created by removing
19 inapplicable rule entries from the intermediate results based on rule dependencies, which are
20 defined by rule matching and suppression settings (e.g. mutex flags and suppression tags). The
21 applicable rules are then processed to calculate the SCI score for that particular source/target pair
22 based on the rule weights. Remediation costs are also calculated based on the cost of
23 updating/upgrading etc. and the mutex settings.

24 [0088] Turning back to Figure 19, the current target is then evaluated against all remaining
25 sources and then the next target is evaluated. As a result, a N x N matrix can be created that
26 shows a compatibility score for each system against each other system. The matrix can be sorted
27 by grouping the most compatible systems. The sorted SCI matrix 60 is comprised of every
28 source/target combination.

1 [0089] Preferably, an SCI report is then generated comprising the SCI matrix 60 (e.g. Figure
2 9) and for each source-target pair details available pertaining to the SCI scoring weights,
3 remediation costs and applicable rules. The details can preferably be pulled for each
4 source/target pair by selecting the appropriate cell.

5 [0090] An example system configuration compatibility analysis is provided below for an
6 arbitrary environment 12 having four servers, namely server A, server B, server C and server D.

7 [0091] The audit engine 46 collects detailed configuration data from the servers 28, which
8 may include, e.g. UNIX™, Linux™, Windows™, AS/400™ etc. The process of collecting the
9 detailed information is herein referred to as an audit. The audit engine 46 collects the data from
10 instrumented candidate systems 12 through various protocols such as simple network
11 management protocol (SNMP), Windows™ management instrumentation (WMI), SSH etc. as
12 shown in Figure 6. Depending on the protocol used and the data to be collected, instrumentation
13 54 on the environment 12 may need to be augmented with additional software such as an agent
14 associated with the analysis program 10.

15 [0092] The web client 34 and web interface allows the user to define the servers 28 to be
16 audited, the actual data to be collected and the audit communication protocol. An example
17 screen shot of an audit request program 110 is shown in Figure 10.

18 [0093] The program 110 enables a user to name and describe the audit so that it can be later
19 identified by entering the requisite information into the entry boxes 112. The audit request
20 information is included in sub-window 114 and the target server information is included in sub-
21 window 122. The type of request, e.g. SNMP is selected using the drop-down box 116. Based
22 on this selection, a list of request parameters are listed in sub-window 118. As shown in the
23 example, the SNMP version and the particular port are listed as well the type of community
24 string. It will be appreciated that any number of request parameters may be listed. Below the
25 request parameters, sub-window 120 provides a visual list of request templates which defines the
26 data that is to be collected and from where it can be obtained. This list can preferably be edited,
27 e.g. by selecting the "Edit" button as shown.

1 [0094] Examples of the request templates that are not necessarily included in Figure 10 are
2 MIB-II, which includes detailed network configuration and statistical data; host resources, which
3 includes system hardware configuration, devices, installed software, installed patches etc.; a
4 configuration monitor, which obtains OS configuration details; a software auditor, which obtains
5 data pertaining to installed patches; a base system workload via a system activity reporting
6 (SAR), which obtains hourly and daily performance statistics of basis properties such as CPU,
7 memory, swap etc.; and extended system workload via SAR, which obtains hourly and daily
8 performance data for additional properties including file I/O, swap I/O, page faults etc. It will be
9 appreciated that any number and form of request templates can be used based on the particular
10 strategies being used for a particular environment 12.

11 [0095] The portion 122 of the audit request program 110 displays a list of the audit targets,
12 which are each server 28 that is to be audited. In the example shown, servers A-D are listed and
13 this list can be edited should the user wish to add or remove certain servers 28. Once the user
14 has selected the desired settings and targets etc., they may choose an "Audit" button to begin the
15 audit, or they may cancel the audit and/or save the audit before it runs.

16 [0096] The audit results are acquired by requesting information specified in the selected
17 audit request templates from the servers 28 in the audited environments 50, and the audited data
18 is stored. The data is then accessed by the analysis program 10 and processed to generate reports
19 and display such reports.

20 [0097] A detailed configuration report 130 is shown in Figure 11. The detailed configuration
21 report organizes the audit data by system and category in the table 132. As shown in Figure 11,
22 both configuration and workload audit data can be viewed by selecting "View". The report 130
23 also displays audit information such as the name for the targeted audit, the type of audit and a
24 timestamp indicating when the audit was performed.

25 [0098] Portions of the detailed information that has been collapsed within the configuration
26 report can be viewed by selecting the appropriate "View" link in the report 130. An example of
27 a detailed configuration information table 136 is also shown in Figure 11. In the example shown,

server information 138 and kernel parameters 136 are displayed with detailed data pertaining to the associated settings, kernels etc.

[0099] The system compatibility analysis, which ultimately generates the SCI matrix 60, compares the detailed configuration settings for the candidate servers 28 with other candidate systems (not shown) based on a pre-defined rule set 43. The resulting analysis yields a system compatibility index score for each candidate server combination. As noted above, these scores are preferably arranged in the graphical SCI 60 format shown in Figure 9 and sorted using a preferred sorting algorithm as discussed above.

[00100] The analysis program 10 supports multiple rule sets to address a variety of server consolidation strategies and scenarios. Example consolidation rule sets include those for UNIX™ systems, Windows™ systems virtualization, SQL™ servers and Oracle™ databases. As discussed, each rule specifies the important configuration property that is to be compared between the candidate servers, the criteria for comparison, a relative weight to compute the SCI score, a rule description, and any rule dependencies. Other fields may also be considered such as the source and/or target instance, conditional flags, mutex etc. An example rule set 76 arranged in a table 146 for an arbitrary system is shown in Figure 12. Each rule 43 and its associated information is listed as a separate entry 148. Various OS settings (not all shown in Figure 12) are assessed including the OS name, OS version, OS kernel bits, memory size, patch level, name service settings, kernel parameters, locale settings, timezone settings etc.

[00101] Total and shared memory criteria tests whether or not the target system has less memory than hypothetical systems that it could be migrated to. As can be seen from Figure 12, a different operating system is deemed to be more important than a different time zone and thus is given a larger weight. The criticality of a rule may vary from system to system and thus the rule sets should be flexible and configurable.

[00102] The SCI score is computed for each candidate system combination by evaluating the rules based on the configuration data obtained from the audit of the environment 12. For each rule that applies to the candidate server pair, the pair's SCI score is reduced by iteratively applying the corresponding rule weight at each iteration (i.e. for each rule), from an initial value

21562811.1

to a final value, the final value being the score. A suitable formula for calculating the SCI score is:

[00103] $SCI_{(n+1)} = SCI_n(1 - Weight);$

[00104] where SCI_n is the current SCI score (initially 100) before the next rule is evaluated, SCI_{n+1} is the new SCI score, and Weight is the rule weight.

[00105] For example, for an arbitrary pair of systems, if the following rules: different operating system (weight=0.5); not running the same kernel bits (weight=0.1); and target host has less memory (weight=0.05); were found to apply to a pair of servers 28, the corresponding SCI score would be computed as follows:

[00106] $SCI_1 = 100(1 - 0.5) = 50$

[00107] $SCI_2 = 50(1 - 0.1) = 45$

[00108] $SCI_3 = 45(1 - 0.05) = 42.75$

[00109] Final SCI score = 42.75.

[00110] As can be seen from the example above, where two servers 28 have different operating systems, different kernel bits and the target has less memory than the baseline data, the configuration compatibility for that server pair is quite low. It is therefore unlikely that these servers could be consolidated since they are not compatible at the configuration level.

[00111] The compatibility program 10 provides a user interface for performing the compatibility analysis as shown in Figure 13. Through such an interface, users can specify the input parameters required to generate the SCI scores.

[00112] In the example shown in Figure 13, the program 150 is used to generate a system compatibility report. An audit that has been saved is loaded into the program 150, which lists identifying information such as the timestamp, number of target servers, and the number of failed

evaluations in the table 152. The user can select the report category from the drop-down list 154, e.g. optimization, inventory, change, compliance, administration etc; and can choose the report type from the drop-down list 156. The report parameters 158 may also be selected which are specific to the type of report being generated and define how the report is structured. The target information 162 is also shown and the target can be removed. Once the user has made the appropriate selections, choosing the “Generate” button creates the SCI matrix 60a shown in Figure 14.

[00113] The SCI matrix 60a is presented as an NxN matrix, where the top of each column indicates the server on which each server listed in the first column is consolidated, and the row name indicates the server being migrated. The SCI 60 shown in Figure 14 provides an example matrix of scores for the example environment 12 including servers A, B, C and D. The SCI scores apply to each possible server consolidation pair. The higher scores indicate more compatibility.

[00114] As noted above, a group sort can be applied to the SCI matrix 60, which includes a difference calculation. The difference calculation between a pair of systems can be illustrated making reference to the SCI matrix 60a shown in Figure 14. From this example, the difference between server A and server B may be computed as follows:

[00115] $\text{Diff} = \sqrt{(\text{square}(S(a,a) - S(a,b)) + \text{square}(S(b,a) - S(b,b)) +$

[00116] $\text{square}(S(c,a) - S(c,b)) + \text{square}(S(d,a) - S(d,b)) +$

[00117] $\text{square}(S(a,a) - S(b,a)) + \text{square}(S(a,b) - S(b,b)) +$

[00118] $\text{square}(S(a,c) - S(b,c)) + \text{square}(S(a,d) - S(b,d)))$

[00119] $= \sqrt{(\text{square}(100 - 90) + \text{square}(90 - 100) +$

[00120] $\text{square}(83 - 79) + \text{square}(52 - 52) +$

[00121] $\text{square}(100 - 90) + \text{square}(90 - 100) +$

1 **[00122]** square(83-79) + square(57-57))

2 **[00123]** where sqrt represents a square root operation.

3 **[00124]** The detailed differences shown in Figures 15 and 16 can be viewed by clicking on the
4 relevant cell 66. As can be seen in the example, migrating server C to server D yields a score of
5 57.

6 **[00125]** Selecting this cell accesses the detailed differences table 178 shown in Figure 16,
7 which shows the important differences between the two systems, the rules and weights that were
8 applied and preferably a remediation cost for making the servers more compatible, the data being
9 collectively referred to by numeral 176. As shown in Figure 15, a summary differences table
10 170 may also be presented when selecting a particular cell, which lists the description of the
11 differences 174 and the weight applied for each difference, to give a high level overview of
12 where the differences arise.

13 **[00126]** The SCI matrix 60 may then be used along with a workload compatibility index
14 (WCI) matrix to perform an overall co-habitation analysis as described in co-pending U.S. Patent
15 Application No. * filed on September 26, 2006, and entitled "System and Method for
16 Determining Compatibility of Computer Systems", the contents of which are incorporated herein
17 by reference.

18 **[00127]** It will be appreciated that rules may instead be used for purposes other than
19 consolidation such as capacity planning, regulatory compliance, change, inventory, optimization,
20 administration etc. and any other purpose where the compatibility and/or differences between
21 systems is useful for analyzing systems 28. It will also be appreciated that the program 10 may
22 also be configured to allow user-entered attributes (e.g. location) that are not available via the
23 auditing process and can factor such attributes into the rules and subsequent analysis.

24 **[00128]** It will further be appreciated that although the examples provided above are in the
25 context of a distributed system of servers, the principles and algorithms discusses are applicable
26 to any system having a plurality of sub-systems where the sub-systems perform similar tasks and
27 thus are capable theoretically of being consolidated. For example, a local network having a

21562811.1

1 number of personal computers (PCs) could also benefit from a consolidation analysis, data in
2 databases could be consolidated etc..

3 **[00129]** Although the invention has been described with reference to certain specific
4 embodiments, various modifications thereof will be apparent to those skilled in the art without
5 departing from the spirit and scope of the invention as outlined in the claims appended hereto.

What is claimed is:

1. A method of evaluating differences between a first data set and a second data set for one or more computer systems comprising:
 - obtaining said first data set and said second data set;
 - selecting a parameter according to a differential rule definition;
 - comparing said parameter in said first data set to said parameter in said second data set;
 - determining if a difference in said parameter exists between said data sets;
 - if said difference exists, applying a weight indicative of the relative importance of said difference in said parameter according to said differential rule definition; and
 - providing an evaluation of said difference according to said weight.
2. A method according to claim 1, said first data set for a baseline computer system and said second data set for a target computer system, said baseline and target computer systems being different systems.
3. A method according to claim 1, wherein said first data set and said second data set are indicative of the state of said same system at different times.
4. A method according to claim 1 comprising comparing a plurality of parameters in said systems, determining if a plurality of corresponding differences exist, and applying a plurality of weights according to a differential rule set comprising a plurality of differential rule definitions, whereby said evaluation provides a score considering said plurality of weights.
5. A method according to claim 4 wherein said score is computed by iteratively decrementing an initial value to produce a final value and applying one of said weights at each iteration, said final value being used as said score.
6. A method according to claim 1 further comprising evaluating said data sets using metadata describing the meaning of said data sets.

7. A method according to claim 1 comprising generating a report including a score for said evaluation and displaying said report on an interface.
8. A method according to claim 1 for evaluating a plurality of different computer systems comprising evaluating each pair of said plurality of computer systems according to said differential rule definition and computing a plurality of scores according to said weight.
9. A method according to claim 8 comprising generating a graphical interface for displaying said plurality of scores, said graphical interface comprising a matrix of cells, each said cell corresponding to a pair of said plurality of computer systems, each row of said matrix indicating one of said plurality of computer systems and each column of said matrix indicating one of said plurality of computer systems, each cell displaying a respective one of said plurality of scores indicating the compatibility of the respective pair of said plurality of systems indicated in the corresponding row and column, and computed according to said weights.
10. A method according to claim 1 wherein said parameter is one of operating system (OS) type, OS version, time zone, total memory, OS patch level, and OS kernel bits.
11. A computer readable differential rule definition for evaluating differences between a first data set and a second data set for one or more computer systems comprising:
 - a parameter for said one or more computer systems; and
 - a weight for said parameter indicative of the importance of a difference in said parameter;wherein said differential rule definition is used by a computer application to perform an evaluation of said difference according to said weight.
12. A rule definition according to claim 11 further comprising a mutex flag indicating whether or not multiple matches of said rule incur a single penalty in said evaluation.

13. A rule definition according to claim 11 further comprising a suppression flag indicating whether or not said rule definition is suppressed according to the processing of another rule definition.
14. A rule definition according to claim 11 further comprising a remediation cost representing the cost of eliminating a deficiency in said configuration parameter in one of said first and second data sets according to the same parameter in the other of said data sets.
15. A rule definition according to claim 11, said first data set for a source system, said rule definition comprising a source field specifying the literal value that must match for said parameter on said source system in order for said rule to match, and a target field specifying the literal value that must match for said parameter in said second data set on a target system in order for said weight to apply.
16. A rule definition according to claim 11 comprising a rule type field indicating whether said differential rule definition applies to audited data directly or to normalized values.
17. A rule definition according to claim 11 comprising a rule description for describing said parameter as it applies to the application of said weight.
18. A rule definition according to claim 11 wherein said parameter is one of operating system (OS) type, OS version, time zone, total memory, OS patch level, and OS kernel bits.
19. A rule set comprising at least one rule definition according to claim 11.

1 ABSTRACT

2
3 A method for evaluating the differences between computer systems is provided. The systems
4 are evaluated using one or more differential rule definitions and one or more differential rule
5 sets. The rules evaluate the differences between specific parameters in the computer systems and
6 are used to generate a score that indicates how different the systems are based on these
7 parameters and preferably what the cost to remedy the difference would be. The rules can be
8 used for consolidation analysis, compliance analysis etc. The rules each include a weight that
9 quantifies the importance of the rule to compatibility and the score is affected accordingly.
10 Systems can be evaluated against each other or against themselves at different instances in time.

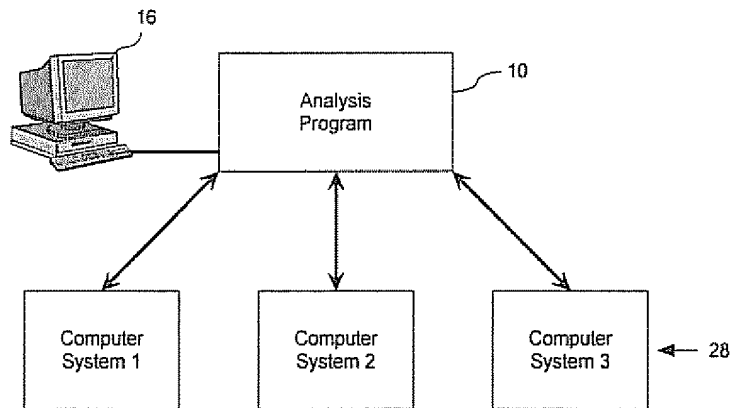


Figure 1

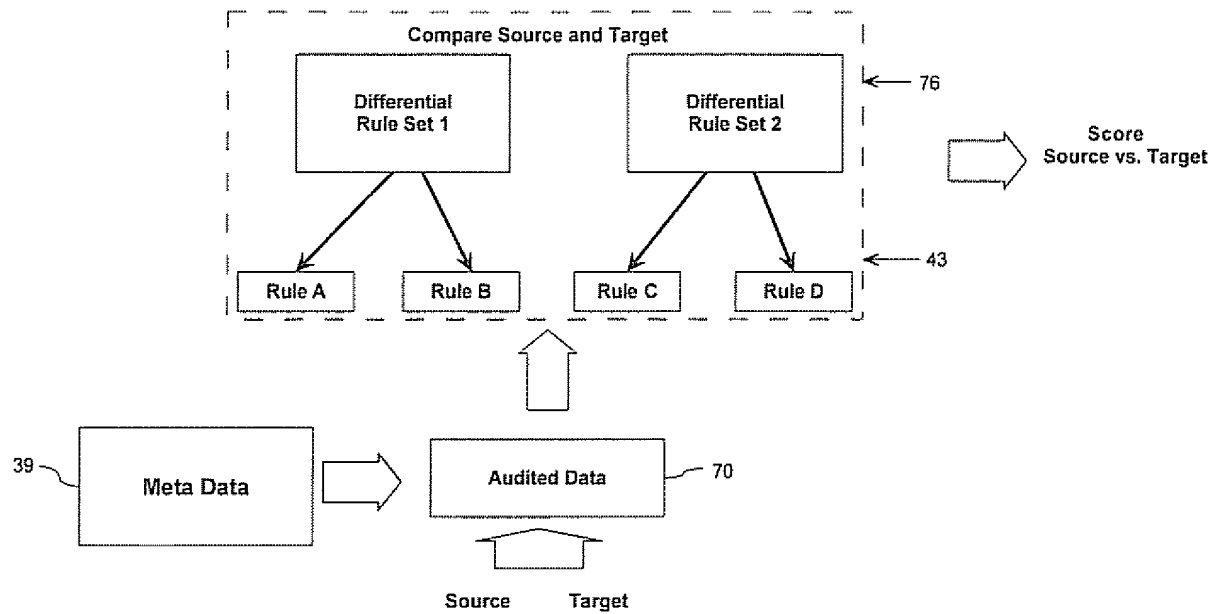


Figure 2




Figure 3

Differential Rule Definition

| Rule Type | Rule Specifier | Source | Target | Weight | Mutex Flag | Match Flag | Supress Flag | Remediation Cost | Description |
|-----------|----------------|--------|--------|--------|------------|------------|--------------|------------------|-------------|
| a | b | c | d | % | Y/N | e | f | \$ | g |

Figure 4

| | Rule Type | Rule Specifier | Baseline | Target | Weight | Mutex Flag | Match Flag | Suppress Flags | Remediation Cost | Description |
|----|------------|--|----------|--------|--------|------------|------------|----------------|------------------|---------------------------------------|
| 1 | AliasQuery | OS | | | 70% | | OS | | \$1,200 | Different Operating Systems |
| 2 | AliasQuery | OS Version | | | 30% | | VER | OS | \$320 | Different Operating System Versions |
| 3 | AliasQuery | Time Zone | | | 20% | | | | \$320 | Systems are in Different Time Zones |
| 4 | AliasQuery | Total memory | | < | 5% | | | | \$350 | Target Has Less Memory |
| 5 | UriQuery | cirba-cm/ServerInformation/OSPatchLevel | | | 5% | | PATCH | OS VER | \$160 | Different Patch Levels |
| 6 | UriQuery | cirba-pa/patchesTable#* | | | 5% | Y | | OS VER | \$80 | Patch Differences Between Systems |
| 7 | UriQuery | cirba-cm/ServerInformation/OSKernelBits | | | 10% | | | OS | \$80 | Not Running Same Kernel Bits |
| 8 | UriQuery | cirba-cm/KernelParametersTable#* | present | absent | 5% | Y | | OS | \$80 | Some Kernel Parameters are Absent |
| 9 | UriQuery | cirba-cm/KernelParametersTable/Setting#SHMMAX | | < | 5% | | | | \$80 | Maximum Shared Memory Settings Differ |
| 10 | UriQuery | cirba-si/applnvTableTable.applnvTableTable/version#oracle | | | 5% | | | | \$3,000 | Different Version of Oracle |
| 11 | UriQuery | cirba-si/applnvTableTable.applnvTableTable/version#oracle | 9 | 8 | 4% | | | | \$1,500 | Different Version of Oracle |
| 12 | UriQuery | cirba-si/applnvTableTable.applnvTableTable/version#apache | | | 2% | | | | \$80 | Different Version of Apache |
| 13 | UriQuery | win32-os/Win32_OperatingSystem/ServicePackMajorVersion | | | 5% | | PATCH | OS VER | \$160 | Different Service Pack Levels |
| 14 | UriQuery | win32-quickfix/win-hotfix#* | present | absent | 5% | Y | | OS VER | \$80 | Patch Differences Between Systems |
| 15 | UriQuery | win32-os/Win32_StartupCommand#* | | | 2% | Y | | OS | \$80 | Startup Commands are Different |
| 16 | UriQuery | win32-misc/Win32_Environment/VariableValue#Path | | | 2% | | | OS | \$80 | Global Path is Different |
| 17 | UriQuery | win32-misc/Win32_Environment/VariableValue#Path <SYSTEM> | | | 2% | | | OS | \$80 | System Path is Different |
| 18 | UriQuery | win32-os/Win32_Service#* | present | absent | 5% | Y | | OS | \$160 | Some Services not On Target |
| 19 | UriQuery | win32-os/Win32_Service/Started | | | 1% | Y | | OS | \$80 | Target |
| 20 | UriQuery | win32-application/win-products#McAfee VirusScan Enterprise | present | absent | 5% | | | OS | \$200 | McAfee Not On Target |

Figure 5

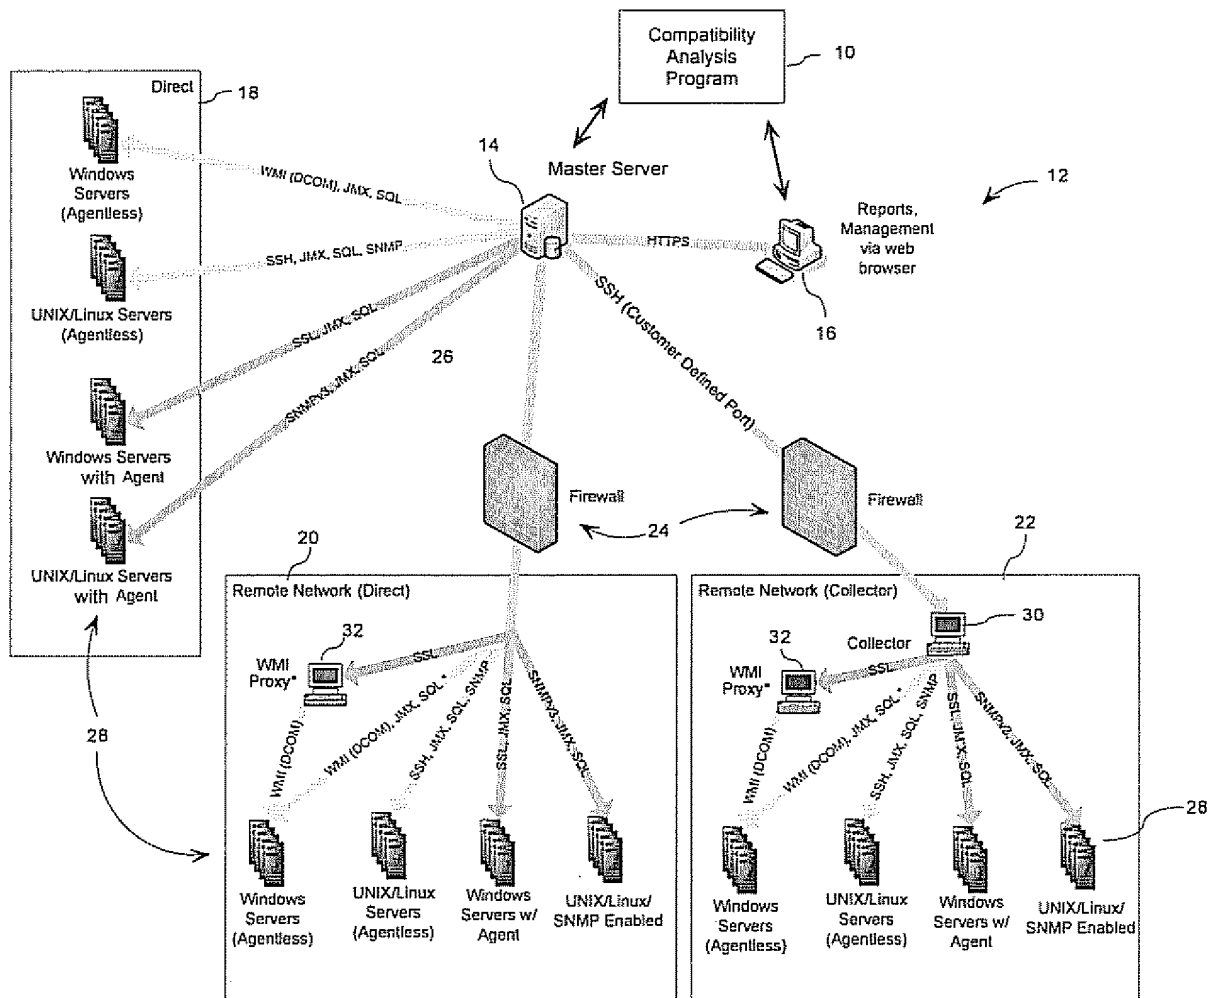


Figure 6

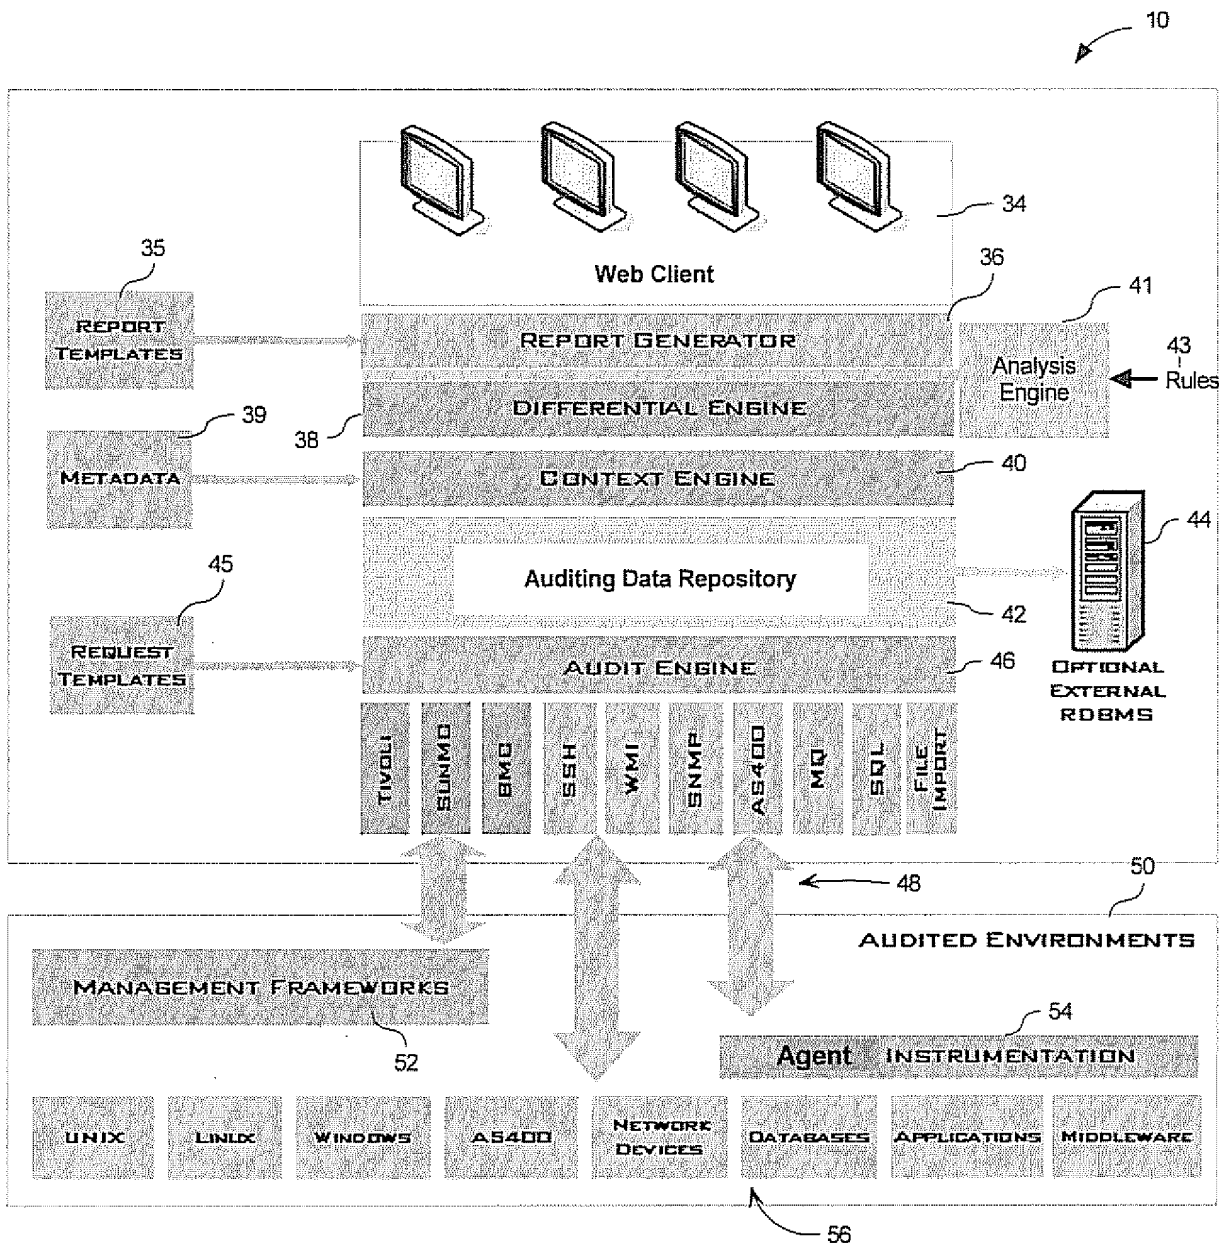


Figure 7

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

102

Consolidation & Virtualization Analysis Strategy Instrumentation Enablement

104

Key:

• Detailed Coverage

◦ Partial Coverage or Requires Customization

105

Database Consolidation

100

Figure 8

☒

AUDIT NAME: COMPANY X } 112

DESCRIPTION: Server Consolidation } 114

REQUEST TYPE: SNMP ☒ ← 116

REQUEST PARAMETERS:

| NAME | Value |
|--------------|------------|
| SNMP Version | 1 |
| Port | 3161 ← 118 |
| Comm. String | public |

REQUEST TEMPLATES: ↙ 120

- /Database_Settings/IETFNetwork_Services
- /Installed_Patches/Patch_Auditor
- /System_Information/Security_Details
- .
- .
- .
- .
- .
- .
- .

EDIT

AUDIT TARGETS: 122

- Server A
- Server B
- Server C
- server D

EDIT

SAVE AUDIT CANCEL } 124

FIGURE 10

✕
DETAILED CONFIGURATION REPORT

CONFIG. BY TYPE
↙ 132

| System | Time | Family | Complete Report | H/W Config. | | BASE WORKLOAD DATA |
|----------|------|--------|-----------------|-------------|------|--------------------|
| SERVER A | : | : | <u>VIEW</u> | : | : | : |
| SERVER B | : | : | : | : | : | : |
| SERVER C | : | : | : | : | : | : |
| SERVER D | : | : | : | : | : | : |

AUDIT INFORMATION
↙ 134

| | NAME | TYPE | TIMESTAMP |
|--------|--------------|---------------|------------------------|
| TARGET | COMPANYX.CH1 | Audit History | MO/DAY/YEAR TIME AM PM |

TOTAL SYSTEMS : 4

↓
↘ 136

SERVER INFORMATION :

| | |
|-----------------|---------|
| O/S Version | 2 |
| No. of Disks | 1024 MB |
| Physical Memory | 502 MHZ |
| CPU Speed | : |
| : | : |
| : | : |
| : | : |

↙ 138

KERNEL PARAMETERS TABLE :

| PARAMETER | VALUE |
|-----------|-----------|
| SYS | SunOS |
| VER | Gen_11822 |
| : | : |
| : | : |
| : | : |
| : | : |

↙ 140

FIGURE 113

146

148

| DATA ELEMENT | CONDITIONAL FLAGS | SOURCE TARGET INSTANCE | WEIGHT | DESCRIPTION | Remediation Cost | MUTEX |
|-----------------|----------------------|------------------------------|--------|----------------------|---------------------|-------|
| SYS_NAME | | | 0.5 | Diff. O/S | — | |
| O/S_Kernel_Bits | | | 0.1 | Not same kernel bits | — | |
| TIME_ZONE | | | 0.02 | Different time zones | — | |
| . | | | . | . | | |
| . | | | . | . | | |
| . | | | . | . | | |
| . | | | . | . | | |
| . | | | . | . | | |
| . | | | . | . | | |
| . | | | . | . | | |
| . | | | . | . | | |

FIGURE 12

AUDIT NAME: COMPANYX.CHI

| TIMESTAMP | TARGETS | FAILED |
|------------------|---------|--------|
| MO/DAY/YEAR TIME | 4 | 0 |

DELETE 152

REPORT CATEGORY: OPTIMIZATION ▾ 154

REPORT TYPE: SYSTEM COMPATIBILITY ▾ 156
WORKLOAD COMPATIBILITY
:

REPORT PARAMETERS:

| NAME | Value |
|--|--|
| ABC123 158 | UNIX ▾ 160 |
| | CPU Utilization → USAGE LIMIT |

TARGETS:

| | | | |
|--|--------------|---------------|------------------|
| REMOVE | CompanyX.CHI | AUDIT History | MO/DAY/YEAR TIME |
|--|--------------|---------------|------------------|

GENERATE 164 162

FIGURE 13

SCI SCORE CARD:

| | SORT ROW | | SORT COLUMN | |
|----------|----------|----------|-------------|----------|
| SCI % | SERVER A | SERVER B | SERVER C | SERVER D |
| SERVER A | 100 | 90 | 83 | 57 |
| SERVER B | 90 | 100 | 79 | 57 |
| SERVER C | 83 | 79 | 100 | 57 |
| SERVER D | 52 | 52 | 52 | 100 |

~60a

FIGURE 14

DIFFERENCES: (MIGRATING A to D)

| DESCRIPTION | WEIGHT |
|------------------------|--------|
| DIFFERENT O/S Versions | 30% |
| DIFFERENT TIME ZONES | 2% |
| ⋮ | ⋮ |
| ⋮ | ⋮ |
| ⋮ | ⋮ |

174 → (points to the table)

172 → (points to 30%)

170 → (points to 2%)

FIGURE 15

DIFFERENCE DETAILS:

| Module | Object | Property | Instance | Baseline | Target | Weight | Remed. Cost |
|-----------------|------------------------------|-----------------|----------|-------------------|---------------------------|--------|-------------|
| Config. Monitor | Name Service Locale settings | Value Time Zone | ipnodes | files US/ Eastern | files.dns Canada/ Eastern | 2% | — |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |

178 → (points to the table)

176 → (points to the table)

FIGURE 16

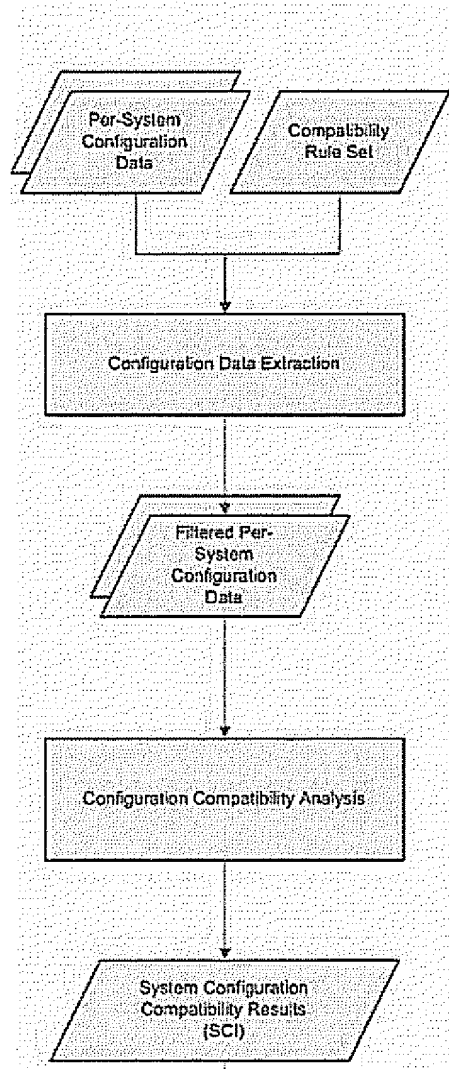


Figure 17

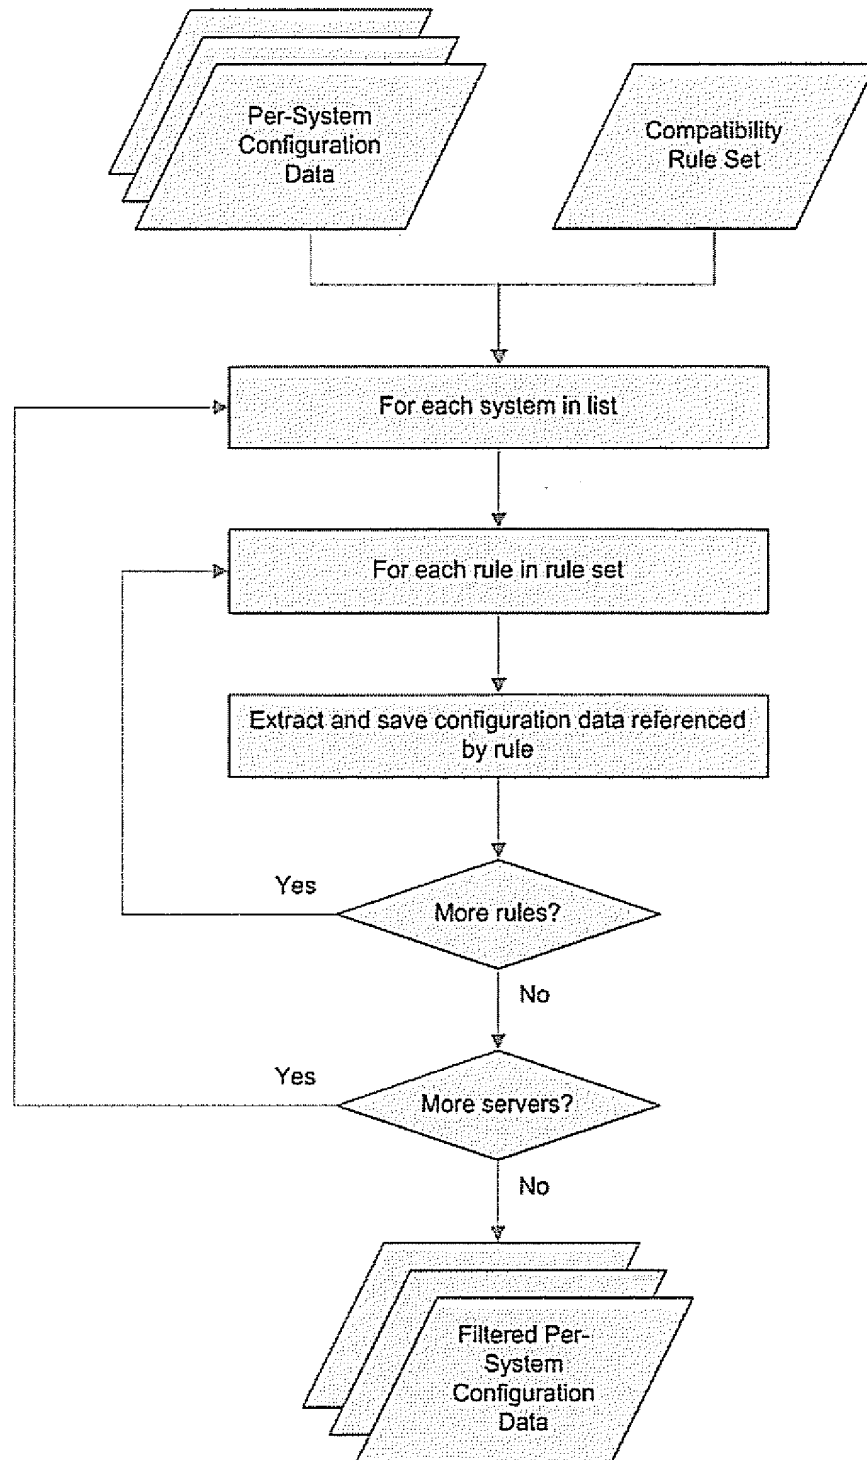


Figure 18

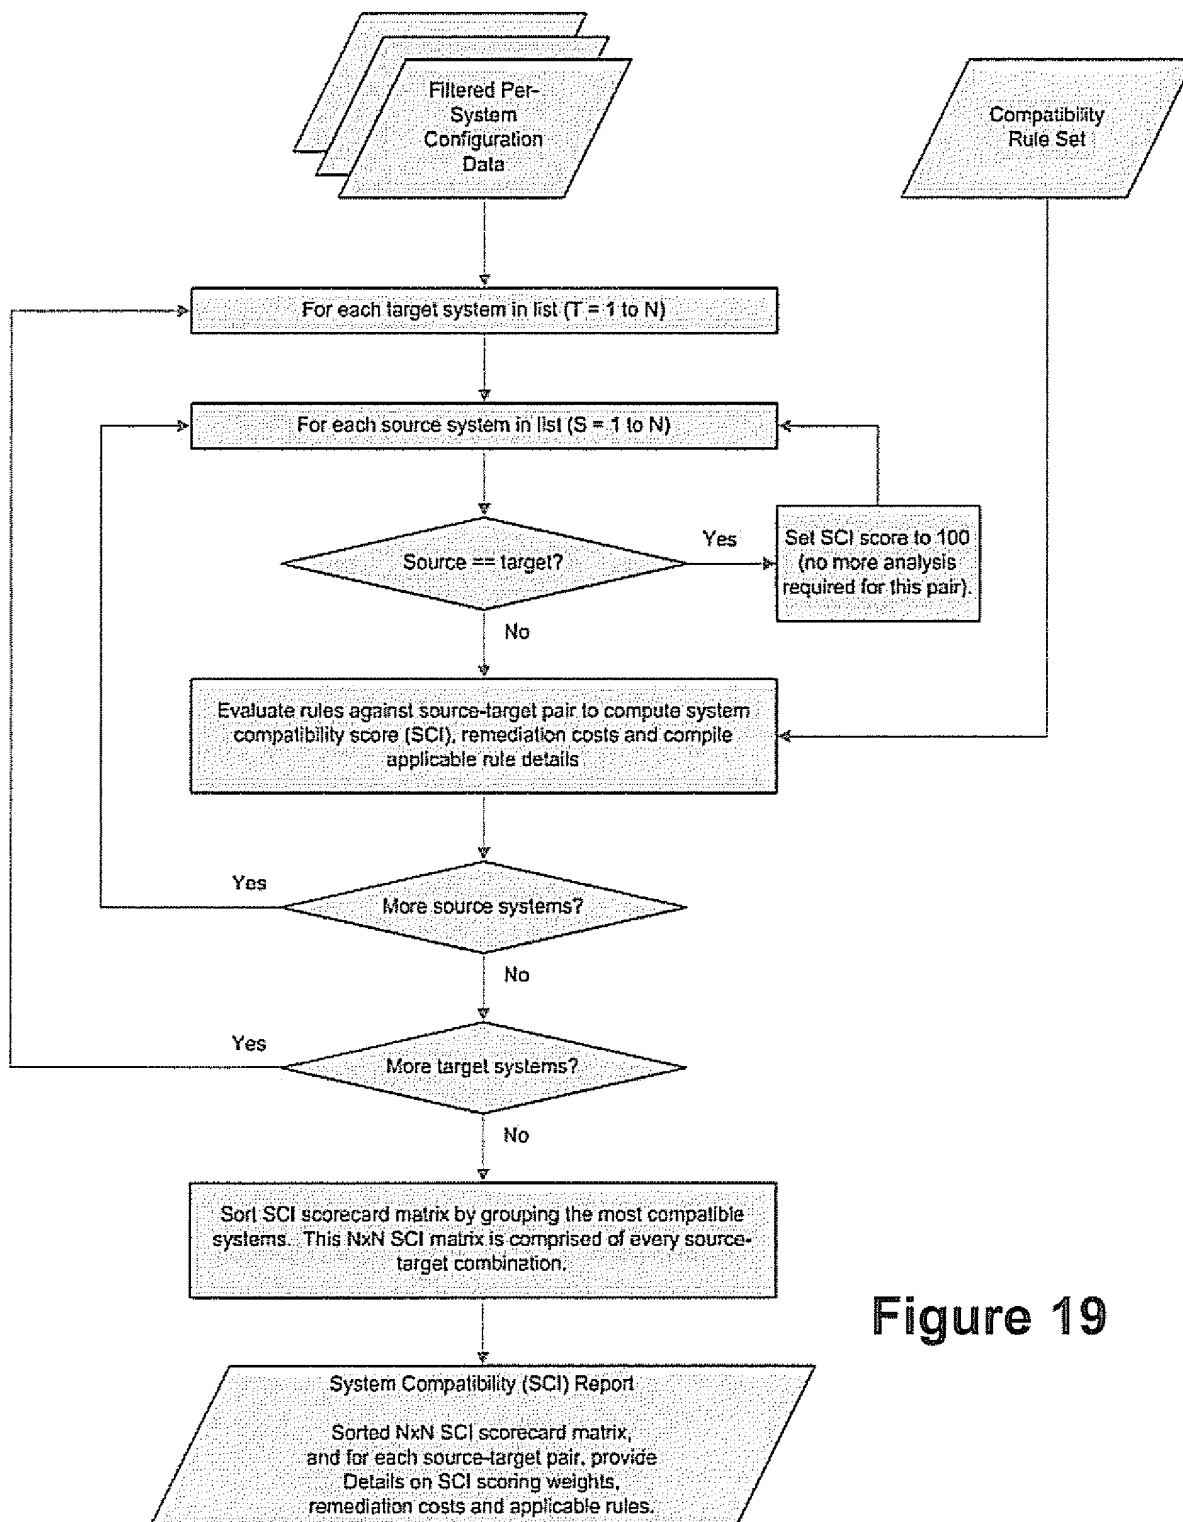


Figure 19

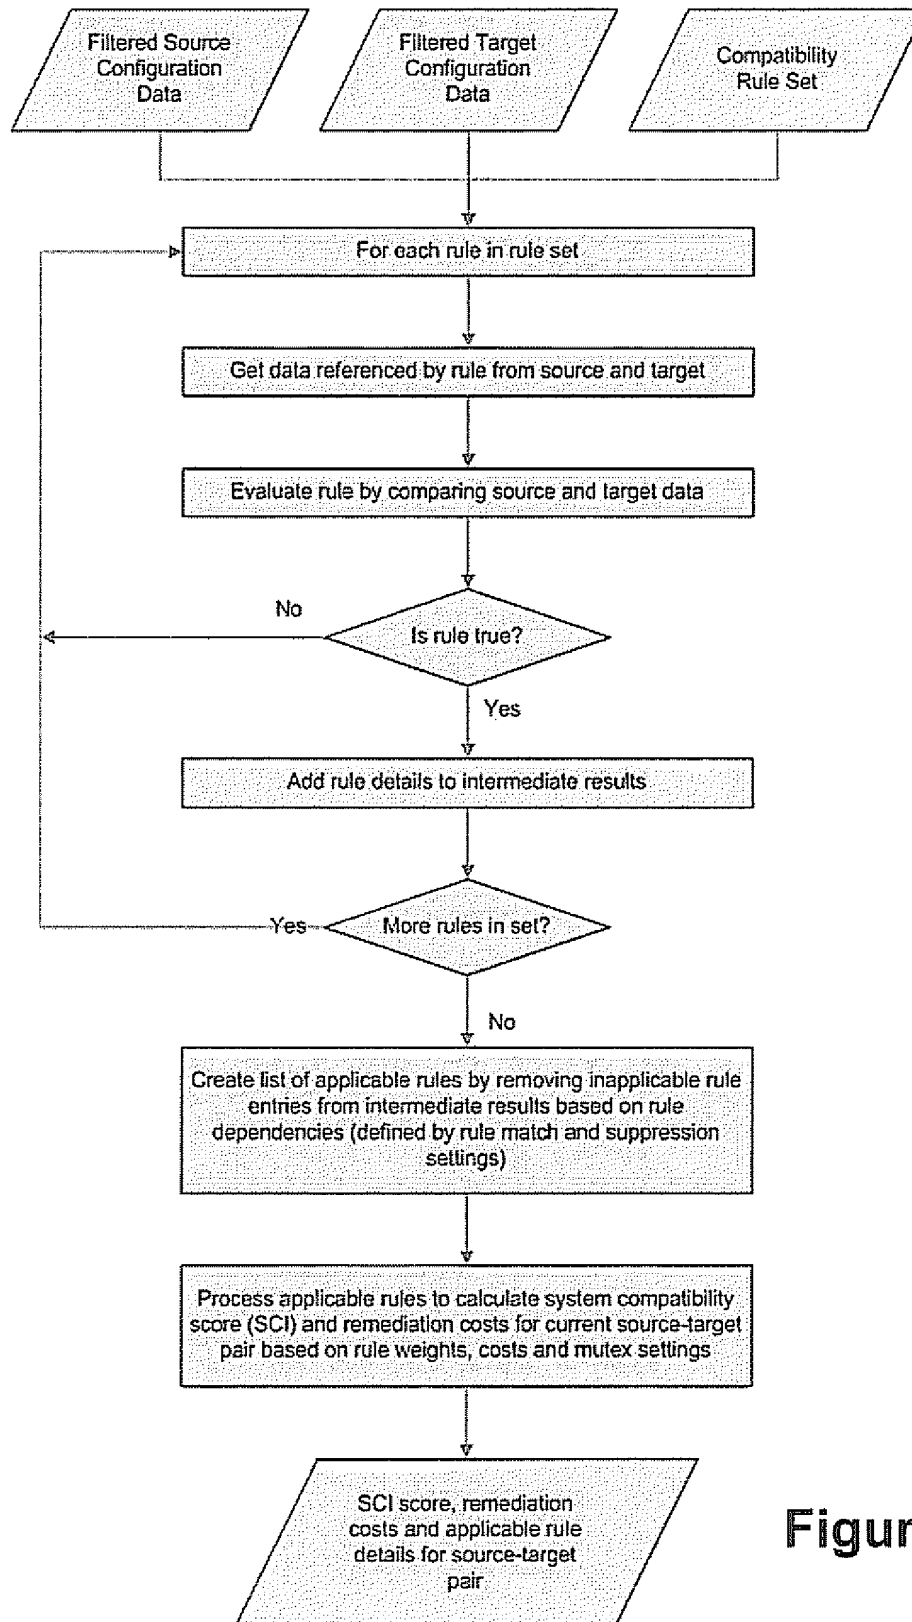


Figure 20