



- (51) International Patent Classification:
H04L 29/06 (2006.01) *H04L 29/08* (2006.01)
- (21) International Application Number:
PCT/US2013/040639
- (22) International Filing Date:
10 May 2013 (10.05.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/645,517 10 May 2012 (10.05.2012) US
- (71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).
- (72) Inventors: JOHNSEN, Bjørn Dag; Vilberggrenda 9, N-0687 Oslo (NO). NARASIMHAMURTHY, Prabhunandan; 500 Oracle Parkway, Redwood Shores, California 94065 (US). MOXNES, Dag Georg; Leirskallbakken 25, N-1164 Oslo (NO). HOLEN, Line; Vitasen 17, N-1900 Fetsund (NO). HODOBA, Predrag; Hoymyrfjellet 10, N-1389 Heggedal (NO).
- (74) Agents: MEYER, Sheldon, R. et al.; Fliesler Meyer LLP, 650 California Street, Fourteenth Floor, San Francisco, California 94108 (US).


(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: SYSTEM AND METHOD FOR SUPPORTING STATE SYNCHRONIZATION IN A NETWORK ENVIRONMENT



(57) Abstract: A system and method can support network management in a network environment. The network environment can include a plurality of configuration daemons (CDs), wherein each CD resides on a switch in the network environment. The CD operates to receive a configuration file that includes a list of known management key (M_Key) values. Furthermore, the CD operates to store the configuration file, and make the configuration file available to a local subnet manager (SM) on the switch, wherein the local SM is associated with a currently used M_Key value. Then, the CD operates to update the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.

WO 2013/170205 A1

SYSTEM AND METHOD FOR SUPPORTING STATE SYNCHRONIZATION IN A NETWORK ENVIRONMENT

COPYRIGHT NOTICE

5 A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise
10 reserves all copyright rights whatsoever.

Field of Invention:

15 **[0001]** The present invention is generally related to computer systems and software, and is particularly related to supporting a network environment.

Background:

[0002] The interconnection network plays a beneficial role in the next generation of super computers, clusters, and data centers. High performance network technology, such as the InfiniBand (IB) technology, is replacing proprietary or low-performance solutions in the high
20 performance computing domain, where high bandwidth and low latency are the key requirements.

[0003] Due to its low latency, high bandwidth, and efficient utilization of host-side processing resources, IB technology has been gaining acceptance within the High Performance Computing (HPC) community as a solution to build large and scalable computer clusters.
25

Summary:

[0004] Described herein are systems and methods for supporting subnet management in a network environment. The network environment can include a plurality of configuration daemons (CDs), wherein each CD resides on a switch in the network environment. The CD operates to
30 receive a configuration file that includes a list of known management key (M_Key) values. Furthermore, the CD operates to store the configuration file, and make the configuration file available to a local subnet manager (SM) on the switch, wherein the local SM is associated with a currently used M_Key value. Then, the CD operates to update the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the
35 network environment.

[0005] Described herein are systems and methods for supporting subnet management in a network environment. A network switch in the network environment can provide a transactional interface, wherein the transactional interface allows a user to interact with the network environment using a transaction. The transactional interface allows a user to group one or more
40 operations in the transaction, and ensures that no conflicting operations are included in the

transaction.

[0006] Described herein are systems and methods for supporting subnet management in a network environment. The network environment can include a plurality of configuration daemons (CDs), wherein a master CD is an active CD on a switch with a master subnet manager(SM).

5 The master CD operates to perform consistency check on one or more states associated with one or more peer CDs in the network environment, and replicate a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

[0007] Also described herein is a system for supporting network management in a network environment. The system comprises means for receiving, via a configuration daemon (CD) on a switch in the network environment, a configuration file, wherein the configuration file includes a list of known management key (M_Key) values. The system further comprises means for storing the configuration file, and means for making the configuration file available to a local subnet manager(SM) on the switch, wherein the local SM is associated with a currently used M_Key value. The system further comprises means for updating the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.

10 **[0008]** Further described herein is a system for supporting network management in a network environment. The system comprises a configuration daemon (CD) on a switch wherein the CD is configured to receive a configuration file that includes a list of known management key (M_Key) values. The system is also configured to store the configuration file, and to make the configuration file available to a local subnet manager(SM) on the switch, wherein the local SM is associated with an currently used M_Key value. The system is configured to update the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.

25 **[0009]** Also described herein is a system for supporting network management in a network environment. The system comprises means for providing a transactional interface at a network switch, wherein the transactional interface allows a user to interact with the network environment using a transaction. The system further comprises means for grouping one or more operations in the transaction and means for ensuring that no conflicting operations are included in the transaction.

30 **[0010]** Further described herein is a system for supporting network management in a network environment. The system comprises a network switch, wherein the network switch is configured to provide a transactional interface, wherein the transactional interface allows a user to interact with the network environment using a transaction. The system is configured to group one or more operations in the transaction and to ensure that no conflicting operations are included in the transaction.

[0011] Also described herein is a system for supporting network management in a network environment. The system comprises means for providing a plurality of configuration daemons (CDs) in the network environment, wherein a master CD is an active CD on a switch with a master subnet manager(SM). The system further comprises means for performing, via the master CD, a consistency check on one or more states associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM and means for replicating a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

[0012] Further described herein is a system for supporting network management in a network environment. The system further comprises a master configuration daemon (CD), wherein the master CD is an active CD on a switch with a master subnet manager(SM), wherein the master CD is configured to be configured to perform a consistency check on one or more states associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM. The system is also configured to replicate a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

Brief Description of the Figures:

[0013] **Figure 1** shows an illustration for supporting a management key (M_Key) protection model in an IB network, in accordance with an embodiment of the invention.

[0014] **Figure 2** shows an illustration for supporting a transactional command line interface (CLI) in a network environment, in accordance with an embodiment of the invention.

[0015] **Figure 3** illustrates an exemplary flow chart for supporting a transactional command line interface (CLI) in a network environment, in accordance with an embodiment of the invention.

[0016] **Figure 4** shows an illustration for supporting subnet management using configuration daemons (CDs) in a network environment, in accordance with an embodiment of the invention.

[0017] **Figure 5** illustrates an exemplary flow chart for supporting subnet management using configuration daemons (CDs) in a network environment, in accordance with an embodiment of the invention.

[0018] **Figure 6** shows an illustration for supporting a management key (M_Key) configuration file in a network environment, in accordance with an embodiment of the invention.

[0019] **Figure 7** shows an illustration for supporting state synchronization between a configuration daemon (CD) and a subnet manager (SM) on a switch in a network environment, in accordance with an embodiment of the invention.

[0020] **Figure 8** illustrates an exemplary flow chart for supporting state synchronization between a configuration daemon (CD) and a subnet manager (SM) in a network environment, in

accordance with an embodiment of the invention.

[0021] **Figure 9** illustrates a functional block diagram to show features in accordance with an embodiment of the invention.

5 **Detailed Description:**

[0022] The invention is illustrated, by way of example and not by way of limitation, in the figures of the accompanying drawings in which like references indicate similar elements. It should be noted that references to “an” or “one” or “some” embodiment(s) in this disclosure are not necessarily to the same embodiment, and such references mean at least one.

10 **[0023]** The description of the invention as following uses the Infiniband (IB) network as an example for a high performance network. It will be apparent to those skilled in the art that other types of high performance networks can be used without limitation.

[0024] Described herein are systems and methods that can support subnet management in a network, such as an IB network.

15

InfiniBand (IB) Network and Management Key (M_Key) protection model

[0025] The IB networks can be referred to as subnets. An IB subnet can include a set of hosts interconnected using switches and point-to-point links. Additionally, the IB subnet can include at least one subnet manager (SM), which is responsible for initializing and bringing up
20 the network, including the configuration of all the switches, routers and host channel adaptors (HCAs) in the subnet.

[0026] **Figure 1** shows an illustration of supporting a management key (M_Key) protection model in an IB network, in accordance with an embodiment of the invention. As shown in Figure 1, a management key, such as an M_Key 102, can be used to protect an IB fabric, such as an IB
25 subnet 100. The values for the M_Key 102 may only be known by fabric administrators 110, which can have administrator access to the switches A-B 103-104 and the designated subnet manager (SM) nodes 101.

[0027] In the IB subnet 100, a secure HCA firmware implementation in HCA 121-124 can keep the type and identity of various fabric nodes well defined. Each of the HCA121-124 can
30 implement a subnet management agent (SMA) component, e.g. SMAs 131-134. Each of the SMAs 131-134 can be associated with an M_Key, e.g. M_Keys 141-144.

[0028] Furthermore, the connected switches A-B 103-104 can be controlled by the fabric administrator 110, which can install new M_Key values 102, out-of-band, on switches 103-104. Thus, any rogue SMA implementation 131-134 may not compromise the fabric administrator 110
35 defined M_Key 102 values that are used in the IB subnet/fabric 100.

[0029] Additionally, the fabric administrator 110 can ensure that there is infinite M_Key 102 lease time on the switches A-B 103-104. Thus, the host based software 161-164, e.g. a host

based subnet manager on different hosts 111-114 (including an operating system 151-154), may not hijack the control of any switch A-B 103-104 in the IB subnet/fabric 100.

[0030] In accordance with an embodiment of the invention, a single M_Key 102 value (or a single set of M_Key values) can be used for various nodes in the in the IB subnet 100 based on the IB specification defined access restrictions. The correct value for a current M_Key 102 may need to be specified before either reading or updating the M_Key 102, since the secure HCA firmware can ensure that the "read protected" M_Key assigned to the local HCA 121-124 is not exposed to local host based software.

[0031] In accordance with an embodiment of the invention, a designated subnet manager 101 can ignore any HCA ports with un-known M_Key value and leave the corresponding link not initialized. The only impact of a hijacked HCA port M_Key can be that the HCA port may not be operational, and the designated subnet manager 101 can prevent host based software from communicating via this HCA port using normal communication, i.e. non-SMP/VL15 based communication.

[0032] Furthermore, the HCA ports may be set up with finite lease time on M_Keys 102, e.g. due to a high availability concern with the subnet manager(s) 101 that maintains the M_Key 102 lease period.

[0033] In accordance with an embodiment of the invention, the M_Key 102 can be created and managed by fabric administrators 110 and stored in secured memory on switches A-B 103-104 and/or HCAs121-124. A microprocessor on a switch A-B 103-104 or a HCA121-124 can access the memory for reading out the M_Key 102 or writing the M_Key 102 into the memory.

Transactional Command Line Interface (CLI)

[0034] In accordance with an embodiment of the invention, a transactional command line interface (CLI) can be operating in a network environment, e.g. from a network switch with the current master subnet manager (SM).

[0035] **Figure 2** shows an illustration of supporting a transactional command line interface (CLI) in a network environment, in accordance with an embodiment of the invention. As shown in Figure 2, a switch 201 in a network environment, e.g. an IB subnet 200, can include a subnet manager (SM) 202 and a configuration daemon (CD) 203.

[0036] Furthermore, the switch 201 can provide a transactional CLI 204, which allows a user to configure and manage the IB subnet 200 using one or more transactions based on M_Key 205. As shown in Figure 2, a transaction 210, which allows the system to perform various operations, can group one or more commands 212 between a start command 111 and a commit command 213.

[0037] In accordance with an embodiment of the invention, a single upgrade transaction, e.g. the transaction 210, may not perform conflicting operations. This can be implemented in a

dynamic way where the initial operation(s) within the transaction 210 can define both the type of transaction 210 and which operations are allowed. For example, the start command 211 to initiate the transaction 210 can include parameter(s) that explicitly defines which operations are allowed.

5 **[0038]** Alternatively, a logical commit operation in a single upgrade transaction 210 can be implemented as several consecutive commit operations, each of which can represent a sub-transaction that groups together a set of similar operations. For example, all remove operations can be implemented as a sub-transaction before all add operations can be implemented as a subsequent sub-transaction. Such an automated scheme can have the advantage of making the
10 normal update scenarios simpler for the user, and may also have the disadvantage of making the handling of a partially failed transaction more complex in terms of defining which operations have been completed and which have not.

[0039] Furthermore, the system allows the user to keep track of operations that can be grouped within a single transaction, in order to simplifying the handling of a partially failed
15 transaction in terms of defining which operations may have been completed and which may have not.

[0040] In accordance with an embodiment of the invention, the transaction start command 211 can ensure that conditions for completing and committing the transaction are fulfilled. If the current conditions are not acceptable, the system allows the user to choose to abort the
20 transaction. Furthermore, the user can re-issue a commit operation in force mode, after the failure of a normal commit due to current conditions.

[0041] Additionally, the command interface can generate a named template 220, which is a configuration that is not committed to the fabric, but instead stored in a local or remote file system. This named template 220 can also be generated from an already committed
25 configuration.

[0042] The command interface allows a named template 220 to be loaded following a transaction start command 211. The named template 220 can also be a parameter for the start command 211. The user is allowed to modify the configuration information from the loaded named template 220 before committing it in the same way as any other configuration update.

30 **[0043]** Furthermore, the system can prepare different named templates 220 for different use cases. Also, based on the named templates 220, the system can perform backup operations to provide protection for the system, after a complete fabric outage, which may cause the loss of configuration information from all peer CD instances 206 in the fabric. Furthermore, the named template 120 allows an administrator to roll back to an earlier consistent configuration, when the
35 administrator has committed an incorrect configuration. Thus, the administrator can avoid performing a potentially very complex sequence of explicit “undo” and “redo” operations.

[0044] In accordance with an embodiment of the invention, by not allowing the operation of

changing current M_Key value 205 and the operation of deleting the (old) current M_Key value 205 in the same transaction, the system can avoid a race condition when a transaction fails into a state where one SM may have a current M_Key that is unknown to one or more of the other SM instances.

5 [0045] **Appendix A** provides further information with respect to providing a transactional command line interface (CLI) in a network environment and various aspects of the platform described throughout this disclosure. The information in **Appendix A** is provided for illustrational purposes and should not be construed to limit all of the embodiments of the invention.

10 [0046] **Figure 3** illustrates an exemplary flow chart for supporting a transactional command line interface (CLI) in a network environment, in accordance with an embodiment of the invention. As shown in Figure 3, at step 301, the system can provide a transactional interface at a network switch, wherein the transactional interface allows a user to interact with the network environment using a transaction. Then, at step 302, the system can group one or more operations in the transaction. Furthermore, at step 303, the system can ensure that no conflicting
15 operations can be included in the transaction.

Configuration Daemon (CD)

[0047] **Figure 4** shows an illustration of supporting subnet management using configuration daemons (CDs) in a network environment, in accordance with an embodiment of the invention.
20 As shown in Figure 4, a network environment, e.g. an IB subnet 400, can include a plurality of network switches, e.g. switches A-E 401-405.

[0048] Each of the network switches A-E 401-405 can include a subnet manager (SM) and a configuration daemon (CD). For example, switch A 401 can include a CD A 421 in addition to a SM A 411. Here, SM A 411 is the master SM that is responsible for managing the subnet 400.
25 Additionally, each of the switches B-E 402-405 includes a standby SM, e.g. one of SMs B-E 412-415, and a CD instance, e.g. one of CDs B-E 422-425.

[0049] Before the system allows an operation to be performed on the IB subnet 400, the master CD A 421 can perform a consistency check for various state information on the peer CD instances B-E 422-425. Such state information can include availability, compatible
30 implementation/protocol version, correct current daemon run-time state, and correct current configuration state/version.

[0050] Then, the master CD A 421 can replicate a new management key (M_Key) configuration file 410 to the peer CDs B-E 422-425, before instructing the CDs B-E 422-425 to update the related SMs B-E 412-415 with a current M_Key value.

35 [0051] In accordance with an embodiment of the invention, the system can ensure that the master SM 411 always have a known M_Key list 430 that is longer than, or at least as long as, any other list maintained by a standby SM instance. Also, the order of replication can ensure that

a new shorter list is applied to the standby SMs before it is applied to the master SM.

[0052] Thus, the master SM may never have a known M_Key list 430 that is shorter than any of the standby SM instances, and the system can use the length of the known M_Key list 430 as a criteria for determining mastership.

5 **[0053]** Furthermore, the order of replication can be based on the type of the transaction. When the transaction is an "addonly" transaction, the new configuration file can be used to update the known lists of the local (master) SM first, followed by the standby SM instances. If the transaction is not "addonly", then the local (master) SM is updated with the new known lists after the new known lists have been replicated to the standby SMs.

10 **[0054]** The replication process is considered successfully completed, after all peer CDs have confirmed receiving and storing the configuration file, and the configuration file have passed the checksum and metadata checks. Then, the master CD A 421 can complete the commit operation by instructing each peer CD B-E 422-425 to update the related SM B-E 412-415 with the new current M_Key value.

15 **[0055]** In accordance with an embodiment of the invention, the master CD A 421 can update the local SM A 411 after all peer CDs B-E 422-425 have been successfully updated. This sequence can be independent of the particular type of a transaction.

[0056] Additionally, the system may not instruct the SM instances to change the current M_Key value until all the SM instances, i.e. the actual SM instances – not just the corresponding
20 CD instances - have received the updated known M_Key list. Thus, a new current M_Key value can always be known to all SMs, even when there is a transaction failure while updating the current value.

[0057] In accordance with an embodiment of the invention, an active CD, e.g. a partition daemon (PD) on an Oracle NM2 network switch, may not accept any commit or start command
25 when the local SM node configuration is not augmented with hostname and internet protocol over InfiniBand (IPOIB) address list. Also, the CDs in the SM node list may need to have an appropriate firmware version that allows a start and/or commit operation.

[0058] Furthermore, the switches in the system (i.e. including the switches in the SM node list) need to have an appropriate firmware version before enabling the use of a secure M_Key
30 setting. Configuring secure M_Key in a system with an incompatible firmware version, e.g. an older NM2 firmware, can cause some local operations, such as enabling/disabling local ports, to become dysfunctional. These dysfunctional local operations can include both the enable/disable operations based on explicit CLI commands, and automated enable/disable operations that are performed during system boot and cable hot-plug.

35 **[0059]** In accordance with an embodiment of the invention, the system can log and report the progress of various operations on the SM nodes. If the commit operation is not started or has not fully completed for all SM nodes, then the relevant information can be logged and displayed

as a warning. Also, when the system uses a "force" option, the operation may be reported as either "fully successful" or "partially successful". On the other hand, when the system does not use the "force" option, the operation may not be reported as "partially successful". Furthermore, if an SM node is lost during a commit operation (i.e. it was present when the commit was commenced) then the operation may report a failure, and the user may then have the option to perform another commit operation, or perform a force commit operation, with the same working configuration.

[0060] Additionally, the system can perform force operations, even when some SM node defined CD instances are not available.

[0061] For example, when an SM node instance was missing at one point in time and then comes back as an incompatible SM node instance, e.g. an older version NM2 node. This system may require the problematic instance to be removed from the SM node list, or be upgraded back to a compatible SM node instance, before any subsequent secret M_Key configuration change can be made.

[0062] In order for a stateful fail-over to work, e.g. with a SM enabled older version NM2, there may not be any use of secret M_Keys in the fabric. The system can perform a *disablesecretmkey* CLI command before the downgrade, or perform a *disablesecretmkey* CLI command with "force" option after the downgrade, as a possible workaround for this case of temporarily downgrading SM node instance.

[0063] Additionally, the *disablesecretmkey* CLI command with "force" option can also be a way to recover a SM node when downgrade already has taken place.

[0064] **Figure 5** illustrates an exemplary flow chart for supporting subnet management using configuration daemons (CDs) in a network environment, in accordance with an embodiment of the invention. As shown in Figure 5, at step 501, the system can provide a plurality of configuration daemons (CDs) in the network environment, wherein a master CD is an active CD on a switch with a master subnet manager(SM). Then, at step 502, the master CD can perform consistency check on one or more states associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM. Furthermore, at step 503, the master CD can replicate a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

The management key (M_Key) configuration file

[0065] In accordance with an embodiment of the invention, a management key (M_Key) configuration file 610 can be transmitted between the configuration daemons (CDs). The management key (M_Key) configuration file can include a trusted management key (M_Key) value that is presented to a user, and one or more key values that are transparently generated based on the trusted M_Key. Additionally, the M_Key configuration file can include one or more

of a fixed function, function parameter, and a locally generated random number.

[0066] Figure 6 shows an illustration of supporting a management key (M_Key) configuration file in a network environment, in accordance with an embodiment of the invention. As shown in Figure 6, a switch 601 in the network environment 600 can include a subnet manager (SM) 602 and a configuration daemon (CD) 603. Furthermore, the switch 601 can provide a command line interface (CLI) 604, e.g. a "smsubnetprotection" CLI interface, to the user.

[0067] While the CLI interface 604 may only present a single set of management key (M_Key) values, e.g. a trusted M_Key 605 value, to the user, the underlying implementation in the switch 601 can use additional sets of keys, which can be included in an M_Key configuration file, e.g. the "PD_M_Key-config.conf" file for NM2 network Switch.

[0068] For example, the M_Key configuration file 610 can include: the trusted M_Keys 611 used for nodes defined as "trusted", the untrusted M_Keys 612 used for nodes that are considered not trusted, and the subnet manager keys (SM_Keys) 613 used for basic SM-SM authentication and negotiation.

[0069] The handling of different M_Keys for trusted and un-trusted nodes can ensure that trusted M_Keys 611 used for trusted nodes (switches in particular) can not be observed/learned by un-trusted nodes. This mechanism can also apply to the cases where the trusted nodes are explicitly authenticated via encryption based protocols whereas the un-trusted nodes are not authenticated at all.

[0070] The M_Key configuration file 610 can contain a list of symbolic and hex M_Key values. Also, the M_Key configuration file 610 can contain an explicit identification of the current M_Key value. In addition to the M_Key information, the M_Key configuration file 610 can also include metadata information including format revision, time stamp and ID (e.g. NM2 hostname + IP) for last update, update revision number and checksum.

[0071] Furthermore, the M_Key configuration file 610 can record each added value as a list of associated values. Also, the CD instance 603 can keep track of the current status for the sets of M_Key values, such as information on all the M_Key values with a specific type. For example, the M_Key configuration file 610 can include multiple lines with three values in each line, where each position in the line defines a key type and a complete line can be defined as "current" or "new-current".

[0072] In accordance with an embodiment of the invention, the management key (M_Key) configuration file 610 can be encrypted using a password 620 (e.g. with a default value). The system can use the encrypted transfer, which is based on a common private password on each node, in order to ensure that the sensitive data is not available to anyone being able to observe data traffic on a management network connecting different CD instances.

[0073] Furthermore, when an "add" command is executed, the trusted M_Key value 605 is

what a user of the CLI 604 relates to. Then, the system can generate the other key values transparently for the CLI user. The additional values can be generated as the sum of the user specified key value and a 64 bit value generated based on a unique permutation of the current password. For example, the at least eight characters in the current password can be permuted
5 in a defined way for each key type - e.g. rotate one, rotate two.

[0074] In accordance with an embodiment of the invention, the system can ensure that any formulas or parameters used to generate additional key values are not part of the information to be replicated and synchronized among the SM nodes. For example, the M_Key values can be generated as part of preparing the update transaction information, without relying on any state
10 for generating keys to be present among the CD or the SM instances when the configuration has been defined.

[0075] Furthermore, new M_Key values can be generated as a function of both a currently defined encryption password and a random value in order to minimize the risk of someone deducting an original password or an original M_Key if either password or original M_Key should
15 become available to someone already knowing the derived M_Key value.

[0076] For example, in addition to the sum of the two 64 bit values, a platform generated random number can be included in the sum, in order to prevent giving away the password value. Thus, a party who knows both an un-trusted M_Key and a trusted M_Key encrypted with the password will not be able to generate the corresponding password or trusted M_Key via a very
20 limited number of arithmetic operations.

[0077] Additionally, the use of incorrect password, e.g. as a result of not updating the password correctly on all SM nodes, can cause a checksum error on the remote side. In order to differentiate the type of checksum error due to incorrect password from the type of checksum error due to file corruption during transfer, there can be a checksum for both the encrypted file
25 and a checksum for the original file version. Here, a checksum error for the encrypted file can indicate a transfer error, whereas a checksum error for the un-encrypted file can indicate an incorrect password. Also, there can be metadata for both the encrypted and un-encrypted files in order to ensure consistency and multiple version compatibility at both levels.

[0078] In accordance with an embodiment of the invention, prior to initiating a transaction,
30 the system can ensure consistent password encryption for each CD.

[0079] A challenge for verifying the password encryption is that the password itself may not be communicated explicitly across the potentially insecure management network. Instead, the system can use the password to encrypt a well known message string (i.e. well known "a priori" by all SM nodes). Then, the verification of consistency of the current password can be achieved
35 by sending the encrypted version of the well known message and having the peer SM nodes to confirm that the decryption of the message results in the original well known message.

[0080] For example, the well known message can be the current password string itself.

Furthermore, a dedicated well known string can be a more robust scheme, since the password string typically has a limited length. Additionally, by maintaining a list of historical password or key values, it is possible for the verification logic to check with older password/key values if the current one fails and thereby identify a case of incomplete updates among the SM nodes. Also, a remote check failure for the current value could lead to encryption and checking of older values to determine whether the remote SM node has yet been updated with the current password value.

[0081] In accordance with an embodiment of the invention, such a verification method can also be used to verify that current public/private key infrastructure is consistent among the involved SM nodes, i.e. to ensure that no accidental update has taken place following authentication/verification during discovery/re-discovery. If public/private key infrastructure is in place, then this can be used for encryption of replicated configuration information in order to avoid maintaining a separate encryption password on each SM node.

Start Logic and Consistency Check

[0082] In accordance with an embodiment of the invention, when a transaction start command is executed from a current master node, the state information on various subnet manager (SM) nodes can be verified so that any inconsistencies in the state information can be detected. Also, a pre-condition for an update transaction can be that the system has a consistent (i.e. commit completed and correct checksums) configuration file with the same configuration update revision number.

[0083] The administrator can ensure that correct configuration can be carried forward in an IB subnet, following a failed update transaction that may cause the state information on the SM nodes to be inconsistent. The administrator can initiate a new update transaction to bring the various SM nodes in synchronization. Additionally, an administrator can verify that the inconsistent state does not represent any unexpected condition, and can perform the diff operations that can explicitly provide the detailed difference in various SM nodes or between a specified pair of SM nodes.

[0084] Furthermore, when the local configuration on the current master SM node is not the most recent one or is incorrect, the system can perform a special operation from the master SM node to fetch the current configuration from another standby SM node, and use this configuration as a starting point for the working configuration when performing a new update transaction. Additionally, a special "force" option can be used to override consistency checks for both start and commit operations, e.g. after the inconsistent state has been analyzed and the correct configuration for synchronizing the various SM nodes have been determined by the administrator.

[0085] In accordance with an embodiment of the invention, when fatal synchronization

problems occurs in the IB subnet, the administrator can be allowed to resolve the underlying problem before continuing or restarting any update transaction. For example, when performing a commit operation, the configuration state of the various SM nodes is expected to be exactly the same as the configuration state when the corresponding start operation was performed, i.e. any
5 change in this period may imply that some unsolicited state change has taken place and may indicate a fatal synchronization error between the various SM nodes. Additionally, at both start and commit time, it can be verified that only a single master is present among the SM nodes. If otherwise, then there is a fatal synchronization problem among the SM nodes.

[0086] Furthermore, the consistency of encryption password value among the various SM
10 nodes can be verified as part of the transaction start operation, and can again be verified in the initial part of the commit operation along with the verification that no configuration change has taken place. Additionally, the state consistency checks can be available as command line interface (CLI) subcommands from any SM node at any point in time.

15 **State Synchronization between the CD and the SM**

[0087] Figure 7 shows an illustration for supporting state synchronization between a configuration daemon (CD) and a subnet manager (SM) on a switch in a network environment, in accordance with an embodiment of the invention. As shown in Figure 7, a switch 701 in an IB subnet 700 can include a subnet manager (SM) 702 and a configuration daemon (CD) 703.

20 **[0088]** When the CD 703 receives a new M_Key configuration file 710 as part of a commit transaction, the CD 703 can store the new M_Key configuration file 710 and make the list of known M_Key values 704 available to the local SM 702. Then, the SM 702 can start using the new list of known M_Key values 704 in subnet probing/discovery operations. The SM 702 may continue using the previously defined current M_Key value 705 when updating/setting M_Key
25 value for any subnet management agent (SMA) instance.

[0089] The CD 703 may not update the current M_Key value 705 used by the SM 702, prior to receiving an instruction from the master CD. Also, the CD 703, when it is a standby CD, can communicate the new list of known values 704 to the local SM 702, before the operation is acknowledged to the master CD, which implements the distributed transaction.

30 **[0090]** Unlike partition configuration updates, the new list of known M_Keys 704 can be made available to the SM 702 immediately, since this does not impose any state change in the subnet, and the manual and automatic procedures can make sure that a new list of known M_Key values 704 may always include any secret M_Key value that may be present in a switch node in the system (i.e. including a switch node with a disabled SM).

35 **[0091]** Furthermore, in order for the list of known M_Key values 704 to be used as part of the consistency check and master election among secret key enabled SMs in the subnet, the order of updates can ensure that the current master always has the largest list of known M_Key

values, and thereby avoid any un-intended master handover in the middle of an update transaction.

[0092] In accordance with an embodiment of the invention, the CD instance 703 can update the local state information about the current M_Key value 705, when the CD instance 703 is instructed by the master CD, which controls the commit operation. Also, the CD instance 703 can instruct the local SM 702 to start using the new current M_Key value, e.g. via a dedicated SM CLI command.

[0093] Furthermore, whenever the SM 702 or CD 703 is restarted, the CD instance 703 can update the SM 702 with the current local secret key configuration. Thus, the system can ensure that the subnet management operation is always consistent with the current configuration, and the system can ensure that no master election is performed without the current secret key configuration in place. Also, the state synchronization between CD 703 and SM 702 during startup can make sure that the SM 702 has the current configuration prior to initiating the initial subnet discovery.

[0094] In accordance with an embodiment of the invention, the CD 703 can use a dedicated SM CLI command to communicate the new list of known M_Key values 704 to the SM 702. For example, the SM 702 may only receive the secret key configuration via the private CLI commands, and may not perform any persistent storing of this configuration (i.e. the secret key configuration may only be part of the SM run-time image).

[0095] The SM CLI command can include a plain list of comma separated hex M_Key values that defines the list of known M_Key values 704. Furthermore, the SM CLI allows three separate lists to be specified, in order to take three independent M_Key values as input, such as the trusted M_Key, the untrusted M_Key and the SM_Key.

[0096] For example, in a NM2 switch, the new current M_Key value can be the same as what was identified as "new current" in the replicated M_Key configuration file. The old current M_Key value can remain to be part of the known M_Key list. (I.e. two update transactions may be needed in order to delete a value that has been used as current.)

[0097] In order to ensure that a new current M_Key value is not activated prematurely, the replicated configuration file may always contain information about the (still) currently used M_Key value as well as the intended new value. When a *set-current* operation takes place (i.e. the corresponding command is received from the master CD implementing the transaction), the new M_Key value can be signaled to the SM 702, and then the recorded currently used M_Key value can be updated to reflect the new specified value before the operation is acknowledged. Thus, the master SM can have a known M_Key list longer than, or at least as long as, any other list maintained by a standby SM.

[0098] Furthermore, the CD instance 703 can keep track of both the actual current M_Key value 705 and the enabled state. For example, the handling of the disabling and enabling of

secret M_Key for the SM 702 can be reflected via a value of zero for the disabled state, and a non-zero current value for the enabled state. Furthermore, the CD instance 703 can provide the value of zero, instead of the current value, to the SM 702, when the system is set to be in a disabled state. Then, the CD instance 703 can signal the (real) current M_Key value to the SM 702, when the system is changed from a disabled state into an enabled state.

[0099] In accordance with an embodiment of the invention, the SM 702 can receive explicit enable/disable notifications for the use of secret M_Keys. Also, the SM can use a readable M_Key as defined in a configuration file, e.g. the opensm.conf file, when the use of secret M_Key is disabled (as a default). Additionally, the readable M_Key value in the configuration file can be maintained by the CLI commands. Furthermore, the M_Key protection bits to be used by the SM 702, which reflects readable M_Key, can be defined in the opensm.conf configuration file. On the other hand, the M_Key protection bits can be implicitly defined when the "secret" mode is enabled via the SM CLI interface.

[00100] Also, the M_Key lease time value can be a constant defined in the configuration file. For example, a constant defined by the NM2 FW revision specific opensm.conf template can be sufficient. Alternatively, the M_Key lease time value can be configurable. This M_Key lease time value can be large (e.g. in the order of at least one minute) so that there is never any practical risk that the SM can not refresh the lease time before the deadline (i.e. as long as at least one SM is sweeping). Additionally, the platform level monitoring daemons can also have a role in refreshing the local lease time.

[00101] **Figure 8** illustrates an exemplary flow chart for supporting state synchronization between a configuration daemon (CD) and a subnet manager (SM) in a network environment, in accordance with an embodiment of the invention. As shown in Figure 8, at step 801, a configuration daemon (CD) on a switch in the network environment can receive a configuration file that includes a list of known management key (M_Key) values. Furthermore, at step 802, the CD can store the configuration file, and make the configuration file available to a local subnet manager(SM) on the switch, wherein the local SM is associated with an currently used M_Key value. Then, at step 803, the CD can update the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment

Failed Update Transaction and Master Subnet Manager (SM) Failover

[00102] When a update transaction is successfully completed, the configuration daemons (CDs) on all subnet manager (SM) nodes can have the same configuration, e.g. with the same configuration update revision number. Additionally, the CDs can either have the same current M_Key value or have no new current M_Key value.

[00103] On the other hand, a failed update transaction may cause an availability or

configuration issue with the current SM nodes in an IB subnet. For example, when an update transaction fails during the replication phase, the CDs may have different current update revision number values (and with different file checksums). Also, if the replication phase is completed successfully, but the update of current M_Key value is not completed, then the CDs may have the same update revision number values and same basic file checksums, but with the different current M_Key value.

[00104] In these scenarios, the administrator can use a CLI command provided by the SM Node switches, e.g. the `listactive` command provided by the NM2 switches, in order to discover and resolve the current status. Then, the administrator can decide how to perform another update transaction on the current master CD, in order to bring the system to a well defined state, e.g. using force options and updating the list of SM nodes in the system.

[00105] If a failed transaction leaves the current persistent configuration with both an old current M_Key value and a new current M_Key value, then any new transaction can start out in this state, and behave as if a `set-current` CLI command had been executed for the M_Key value that is reflected as "new" in the current configuration file. Furthermore, the user can be aware of this current status, and is allowed to modify the setting if needed. Additionally, a failure to observe the rules for what operations can be performed following a failed transaction can lead to an incorrect master election in the subnet, e.g. causing either multiple or no master SMs in the subnet.

[00106] In accordance with an embodiment of the invention, the transaction start command can present all relevant state information for the configuration on all SM nodes to the user. Also, the transaction start command can make sure that the user is aware of any inconsistencies, e.g. after a failure of an updated transaction. Furthermore, the transaction start command can perform special force actions in order to initiate a transaction, and allows the user to select the starting point configuration from any of the available SM nodes.

[00107] If the active master is operational after a failure of an updated transaction, the first step to handle the failed update transaction is to check availability and configuration of other SM nodes, potentially based on error logs/messages from the failed update transaction. If the problem can be corrected, then the same transaction can be retried (i.e. without use of force mode). If the problem can not be corrected (e.g. an SM node - potentially the old master - is no longer operational/available), then any delete operation may not be re-tried in force mode, since this operation may delete a key value that may still be in use by an enabled SM.

[00108] In accordance with an embodiment of the invention, a disable operation can be retried in force mode, since the disable operation does not change the current known key list and does not impose any danger of incorrect master election procedures. Furthermore, depending upon the SM priorities and globally unique identifiers (GUIDs) of the currently available SM nodes relative to the failed SM node(s), the disabled or enabled state of the system may change

as the failed SM nodes become operational again.

[00109] Furthermore, force retrying a failed enable operation is safe, since it does not impose any danger of incorrect master election procedures. Also, the state of the system in terms of enabled/disabled state of the master may not be certain when the same key configuration is present.

[00110] Additionally, force retrying an add operation is safe, since it may increase the number of known keys and thereby it can ensure that the SM nodes have received this update and may take precedence in terms of master election relative to the currently unavailable SM node(s).

[00111] In accordance with an embodiment of the invention, by setting a newly added key as current, the system can be configured for fencing off a SM on the unavailable SM node(s). Thus, the SM on the unavailable SM node(s) may not influence the state of any node (potentially including itself) that has received the updated secret M_Key value.

[00112] The system can use force-add and set-current operations to ensure the fencing off an SM node with un-defined state, without any need for physical service actions. Such fencing mechanism allows the SM nodes configuration to be changed, and also allows the current partition policy to be updated.

[00113] When a leaf switch is in an undefined state that can cause the partitioning state of the directly connected hosts to be not controllable, the partitioning configuration may not be updated or changed in a way that conflicts with the current configuration that is assumed to be implemented by an active SM on that leaf switch.

[00114] In order to update or change the partitioning configuration, the leaf switch can be brought to a well defined state with no active SM, or the SM can be brought in synchronization with the rest of the SM nodes, via some additional service actions.

[00115] In accordance with an embodiment of the invention, the system can support master SM failover in the middle of an M_Key update transaction, which follows the same ACID transaction principles as partition policy updates.

[00116] The system can distribute the new list of known M_Key values first, and define a new M_Key value as current for all the SMs in the final part of the commit operation. Also, the enhanced master election protocol can ensure that the elected master can always be the SM that has the longest list of known key values and can be best suited to handle the current configuration.

[00117] Thus, the system can ensure that there may never be a case where a new master SM is not be able to discover/configure the subnet because of lack of knowledge about the most recent M_Key value. For example, all SM enabled switch nodes can have the same new policy if the policy is enabled on any one of switch nodes, otherwise if the policy has not been replicated to all SMs, then no SM can have the new policy enabled.

[00118] **Figure 9** illustrates a functional block diagram to show features of an embodiment of

the invention. The present features may be implemented as system 900 for supporting network management in a network environment. The system 900 includes one or more microprocessors and a configuration daemon (CD) 910 on a switch running on the one or more microprocessors. The CD 910 operates to receive at the receiving unit 920 a configuration file that includes a list of
5 known management key (M_Key) values; to store in the storing unit 930 the configuration file, to make the configuration file available to a local subnet manager (SM) on the switch, wherein the local SM is associated with an currently used M_Key value; and to update on the updating unit 940, the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.

10 **[00119]** According to one embodiment, there is disclosed a system for supporting network management in a network environment on one or more microprocessors. The system comprises means for receiving, via a configuration daemon (CD) on a switch in the network environment, a configuration file, wherein the configuration file includes a list of known management key (M_Key) values. The system also comprises means for storing the configuration file, and
15 making the configuration file available to a local subnet manager(SM) on the switch, wherein the local SM is associated with a currently used M_Key value. The system further comprises means for updating the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.

[00120] Preferably the system comprises means for including in the configuration file a trusted
20 M_Key and one or more key values that are transparently generated based on the trusted M_Key.

[00121] Preferably the system comprises means for continuing using the old current M_Key when updating/setting M_Key value for any subnet management agent (SMA) instance until the local SM is updated with the new M_Key.

25 **[00122]** Preferably the system comprises means for using, via the CD, a dedicated SM command line interface (CLI) command to update the local SM.

[00123] Preferably the system comprises means for updating, via the CD, the local SM with current local secret key configuration when at least one of the SM and CD is restarted.

[00124] Preferably the system comprises means for containing in the configuration file
30 information about the currently used M_Key value as well as the intended new M_Key value.

[00125] Preferably the system comprises means for recording the currently used M_Key value in the list of known M_Key values.

[00126] Preferably the system comprises means for using the new M_Key value to fence off a SM on an unavailable node.

35 **[00127]** Preferably the system comprises means for ensuring that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.

[00128] Preferably the system comprises means for initiating a commit operation on the

master CD.

[00129] According to one embodiment, there is disclosed a system for supporting network management in a network environment. The system comprises a configuration daemon (CD) on a switch wherein the CD is configured to receive a configuration file that includes a list of known management key (M_Key) values. The system is configured to store the configuration file, and to make the configuration file available to a local subnet manager(SM) on the switch, wherein the local SM is associated with a currently used M_Key value. The system is also configured to update the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.

[00130] Preferably the system is capable of having said configuration file include a trusted M_Key and one or more key values that are transparently generated based on the trusted M_Key.

[00131] Preferably the system is capable of having said local SM continue using the old current M_Key when updating/setting M_Key value for any subnet management agent (SMA) instance until the local SM is updated with the new M_Key.

[00132] Preferably the system is capable of having said CD operate to use a dedicated SM command line interface (CLI) command to update the local SM.

[00133] Preferably the system is capable of having said CD operate to update the local SM with current local secret key configuration when at least one of the SM and CD is restarted.

[00134] Preferably the system is capable of having said configuration file contain information about the currently used M_Key value as well as the intended new M_Key value.

[00135] Preferably the system is capable of having said CD operate to record the currently used M_Key value in the list of known M_Key values.

[00136] Preferably the system is capable of having said new M_Key value used to fence off a SM on an unavailable node.

[00137] Preferably the system is capable of having said master CD operate to ensure that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.

[00138] According to one embodiment, there is disclosed a system comprising supporting network management in a network environment on one or more microprocessors. The system comprises means for providing a transactional interface at a network switch, wherein the transactional interface allows a user to interact with the network environment using a transaction. The system further comprises means for grouping one or more operations in the transaction and means for ensuring that no conflicting operations are included in the transaction.

[00139] Preferably the system comprises means for associating the network switch with a subnet manager (SM) in the network environment.

[00140] Preferably the system comprises means for allowing the transactional interface to be

a command line interface (CLI).

[00141] Preferably the system comprises means for using one or more commands to perform the one or more operations, wherein the one or more commands are grouped between a start command and a commit command in the transaction.

5 **[00142]** Preferably the system comprises means for ensuring that one or more conditions for completing and committing the transaction are fulfilled at the beginning of the transaction.

[00143] Preferably the system comprises means for allowing the user to abort the transaction if one or more current conditions are not acceptable.

[00144] Preferably the system comprises means for allowing the user to issue a commit
10 operation in force mode, in which case the commit operation can succeed when one or more configuration daemons are not available.

[00145] Preferably the system comprises means for initiating an update transaction following a previously failed update transaction to bring a plurality of SM nodes in synchronized.

[00146] Preferably the system comprises means for enabling a secret management key
15 (M_Key) before starting the transaction.

[00147] Preferably the system comprises means for setting an defined M_Key value as a M_Key value that will be used by a SM following a subsequent commit operation.

[00148] According to one embodiment, there is disclosed a system for supporting network
20 management in a network environment. The system comprises a network switch, the wherein the network switch is configured to provide a transactional interface, wherein the transactional interface allows a user to interact with the network environment using a transaction. The system is further configured to group one or more operations in the transaction and to ensure that no conflicting operations are included in the transaction.

[00149] Preferably the system is capable of having a network switch associated with a subnet
25 manager (SM) in the network environment.

[00150] Preferably the system is capable of having the transactional interface a command line interface (CLI).

[00151] Preferably the system is capable of having the network switch operate to use one or
30 more commands to perform the one or more operations, wherein the one or more commands are grouped between a start command and a commit command in the transaction.

[00152] Preferably the system is capable of having the network switch operate to ensure that one or more conditions for completing and committing the transaction are fulfilled at the beginning of the transaction.

[00153] Preferably the system is capable of having the network switch operate to allow the
35 user to abort the transaction if one or more current conditions are not acceptable.

[00154] Preferably the system is capable of having the network switch operate to allow the user to issue a commit operation in force mode, in which case the commit operation can

succeed when one or more configuration daemons are not available.

[00155] Preferably the system is capable of having the network switch operate to allow the user to initiate an update transaction following a previously failed update transaction to bring a plurality of SM nodes in synchronized.

5 **[00156]** Preferably the system is capable of having the network switch operate to allow the user to enable a secret management key (M_Key) before starting the transaction.

[00157] According to one embodiment, there is disclosed a system for supporting network management in a network environment on one or more microprocessors. The system comprises means for providing a plurality of configuration daemons (CDs) in the network environment, wherein a master CD is an active CD on a switch with a master subnet manager(SM). The system also comprises means for performing, via the master CD, consistency check on one or more states associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM. The system further comprises means for replicating a configuration file to the one or more peer CDs, wherein the configuration
10 file includes at least one management key (M_Key) value.

[00158] Preferably the system comprises means for instructing, via the master CD, the one or more peer CDs to update the standby SMs with the at least one M_Key value.

[00159] Preferably the system comprises means for determining an order for replicating the configuration file to the one or more peer CDs based on a type of a transaction.

20 **[00160]** Preferably the system comprises means for using the configuration file to first update one or more known lists of M_Key values on the master SM, when the transaction is an add-only transaction.

[00161] Preferably the system comprises means for using the M_Key configuration file to first update one or more known lists of M_Key values on the one or more standby SMs, when the
25 transaction is not an add-only transaction.

[00162] Preferably the system comprises means for allowing the one or more states to include availability of the one or more peer CDs, compatible implementation/protocol version, correct current daemon run-time state, and correct current configuration state/version on the one or more peer CDs.

30 **[00163]** Preferably the system comprises means for verifying consistent encryption password for each CD prior to initiating a transaction.

[00164] Preferably the system comprises means for prohibiting changing a new current M_Key value and deleting an old current M_Key value in a same transaction.

[00165] Preferably the system comprises means for ensuring that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.
35

[00166] Preferably the system comprises means for initiating a commit operation on the master CD.

[00167] According to one embodiment, there is disclosed a system for supporting network management in a network environment. The system comprises a master configuration daemon (CD), wherein the master CD is an active CD on a switch with a master subnet manager (SM), wherein the master CD is configured to perform consistency check on one or more states
5 associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM. The system is also configured to replicate a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

[00168] Preferably the system is capable of having the master CD operate to instruct the one
10 or more peer CDs to update the standby SMs with the at least one M_Key value.

[00169] Preferably the system is capable of having an order for replicating the configuration file to the one or more peer CDs be determined based on a type of a transaction.

[00170] Preferably the system is capable of having the master CD operate to use the configuration file to first update one or more known lists of M_Key values on the master SM,
15 when the transaction is an add-only transaction.

[00171] Preferably the system is capable of having the master CD operate to use the M_Key configuration file to first update one or more known lists of M_Key values on the one or more standby SMs, when the transaction is not an add-only transaction.

[00172] Preferably the system is capable of having the one or more states include availability
20 of the one or more peer CDs, compatible implementation/protocol version, correct current daemon run-time state, and correct current configuration state/version on the one or more peer CDs.

[00173] Preferably the system is capable of having a consistent encryption password verified for each CD prior to initiating a transaction.

[00174] Preferably the system is capable of having an operation to change a new current
25 M_Key value and an operation to delete an old current M_Key value is prohibited in a same transaction.

[00175] Preferably the system is capable of having the master CD operate to ensure that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM
30 instance.

[00176] According to an embodiment, there is disclosed a method for supporting network management in a network environment, comprising providing a transactional interface at a network switch, wherein the transactional interface allows a user to interact with the network environment using a transaction; grouping one or more operations in the transaction; and
35 ensuring that no conflicting operations are included in the transaction.

[00177] According to an embodiment, there is disclosed a method for associating the network switch with a subnet manager (SM) in the network environment.

[00178] According to an embodiment, there is disclosed a method that is capable of allowing the transactional interface to be a command line interface (CLI).

[00179] According to an embodiment, there is disclosed a method for using one or more commands to perform the one or more operations, wherein the one or more commands are grouped between a start command and a commit command in the transaction.

[00180] According to an embodiment, there is disclosed a method for ensuring that one or more conditions for completing and committing the transaction are fulfilled at the beginning of the transaction.

[00181] According to an embodiment, there is disclosed a method for allowing the user to abort the transaction if one or more current conditions are not acceptable.

[00182] According to an embodiment, there is disclosed a method for allowing the user to issue a commit operation in force mode, in which case the commit operation can succeed when one or more configuration daemons are not available.

[00183] According to an embodiment, there is disclosed a method of initiating an update transaction following a previously failed update transaction to bring a plurality of SM nodes in synchronized.

[00184] According to an embodiment, there is disclosed a method for enabling a secret management key (M_Key) before starting the transaction.

[00185] According to an embodiment, there is disclosed a method for setting an defined M_Key value as a M_Key value that will be used by a SM following a subsequent commit operation.

[00186] According to an embodiment, there is disclosed a system for supporting network management in a network environment, comprising one or more microprocessors; a network switch, running on the one or more microprocessors, wherein the network switch operates to provide a transactional interface, wherein the transactional interface allows a user to interact with the network environment using a transaction; group one or more operations in the transaction; and ensure that no conflicting operations are included in the transaction.

[00187] According to an embodiment, there is disclosed a system, wherein the network switch is associated with a subnet manager (SM) in the network environment.

[00188] According to an embodiment, there is disclosed a system, wherein the transactional interface is a command line interface (CLI).

[00189] According to an embodiment, there is disclosed a system wherein the network switch operates to use one or more commands to perform the one or more operations, wherein the one or more commands are grouped between a start command and a commit command in the transaction.

[00190] According to an embodiment, there is disclosed a system wherein the network switch operates to ensure that one or more conditions for completing and committing the transaction

are fulfilled at the beginning of the transaction.

[00191] According to an embodiment, there is disclosed a system wherein the network switch operates to allow the user to abort the transaction if one or more current conditions are not acceptable.

5 **[00192]** According to an embodiment, there is disclosed a system wherein the network switch operates to allow the user to issue a commit operation in force mode, in which case the commit operation can succeed when one or more configuration daemons are not available.

[00193] According to an embodiment, there is disclosed a system wherein the network switch operates to allow the user to initiate an update transaction following a previously failed update
10 transaction to bring a plurality of SM nodes in synchronized.

[00194] According to an embodiment, there is disclosed a system wherein the network switch operates to allow the user to enable a secret management key (M_Key) before starting the transaction.

[00195] According to an embodiment, there is disclosed a non-transitory machine readable
15 storage medium having instructions stored thereon that when executed cause a system to perform the steps comprising providing a transactional interface at a network switch, wherein the transactional interface allows a user to interact with a network environment using a transaction; grouping one or more operations in the transaction; and ensuring that no conflicting operations are included in the transaction.

20 **[00196]** According to an embodiment, there is disclosed a method for supporting network management in a network environment, comprising providing a plurality of configuration daemons (CDs) in the network environment, wherein a master CD is an active CD on a switch with a master subnet manager(SM); performing, via the master CD, consistency check on one or more states associated with one or more peer CDs in the network environment, wherein each
25 said peer CD is an active CD on a switch with a standby SM; and replicating a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

[00197] According to an embodiment, there is disclosed a method for instructing, via the master CD, the one or more peer CDs to update the standby SMs with the at least one M_Key
30 value.

[00198] According to an embodiment, there is disclosed a method for determining an order for replicating the configuration file to the one or more peer CDs based on a type of a transaction.

[00199] According to an embodiment, there is disclosed a method for using the configuration file to first update one or more known lists of M_Key values on the master SM, when the
35 transaction is an add-only transaction.

[00200] According to an embodiment, there is disclosed a method for using the M_Key configuration file to first update one or more known lists of M_Key values on the one or more

standby SMs, when the transaction is not an add-only transaction.

[00201] According to an embodiment, there is disclosed a method for allowing the one or more states to include availability of the one or more peer CDs, compatible implementation/protocol version, correct current daemon run-time state, and correct current configuration state/version on the one or more peer CDs.

[00202] According to an embodiment, there is disclosed a method for verifying consistent encryption password for each CD prior to initiating a transaction.

[00203] According to an embodiment, there is disclosed a method for prohibiting changing a new current M_Key value and deleting an old current M_Key value in a same transaction.

[00204] According to an embodiment, there is disclosed a method for ensuring that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.

[00205] According to an embodiment, there is disclosed a method for initiating a commit operation on the master CD.

[00206] According to an embodiment, there is disclosed a system for supporting network management in a network environment, comprising one or more microprocessors; a master configuration daemon (CD) running on the one or more microprocessors, wherein the master CD is an active CD on a switch with a master subnet manager(SM), wherein the master CD operates to perform consistency check on one or more states associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM; and replicate a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

[00207] According to an embodiment, there is disclosed a system wherein the master CD operates to instruct the one or more peer CDs to update the standby SMs with the at least one M_Key value.

[00208] According to an embodiment, there is disclosed a system wherein an order for replicating the configuration file to the one or more peer CDs can be determined based on a type of a transaction.

[00209] According to an embodiment, there is disclosed a system wherein the master CD operates to use the configuration file to first update one or more known lists of M_Key values on the master SM, when the transaction is an add-only transaction.

[00210] According to an embodiment, there is disclosed a system wherein the master CD operates to use the M_Key configuration file to first update one or more known lists of M_Key values on the one or more standby SMs, when the transaction is not an add-only transaction.

[00211] According to an embodiment, there is disclosed a system wherein the one or more states include availability of the one or more peer CDs, compatible implementation/protocol version, correct current daemon run-time state, and correct current configuration state/version on

the one or more peer CDs.

[00212] According to an embodiment, there is disclosed a system wherein consistent encryption password is verified for each CD prior to initiating a transaction.

5 **[00213]** According to an embodiment, there is disclosed a system wherein an operation to change a new current M_Key value and an operation to delete an old current M_Key value is prohibited in a same transaction.

[00214] According to an embodiment, there is disclosed a system wherein the master CD operates to ensure that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.

10 **[00215]** According to an embodiment, there is disclosed a non-transitory machine readable storage medium having instructions stored thereon that when executed cause a system to perform the steps comprising providing a plurality of configuration daemons (CDs) in a network environment, wherein a master CD is an active CD on a switch with a master subnet manager(SM); performing, via the master CD, consistency check on one or more states
15 associated with one or more peer CDs in the network environment, wherein each said peer CD is an active CD on a switch with a standby SM; and replicating a configuration file to the one or more peer CDs, wherein the configuration file includes at least one management key (M_Key) value.

20 **[00216]** The present invention may be conveniently implemented using one or more conventional general purpose or specialized digital computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art.

25 **[00217]** In some embodiments, the present invention includes a computer program product which is a storage medium or computer readable medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks,
30 ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

35 **[00218]** The foregoing description of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled

in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalence.

Appendix A

A sample CLI Interface

5 **[00219]** For example, the Oracle NW2 network switch can provide a smsubnetprotection CLI interface to support the transactional protocol as described in the above sections. The smsubnetprotection CLI interface can provide the following commands:

start

10 **[00220]** The *start* command can check (via CD) whether all required CD instances ("smnodes") are available and operational. If so, the system can establish a working copy of the current M_Key configuration. The *start* command can include options to force start, to implicitly enable, and to specify a type (e.g. "addonly" or "deleteonly") for the transaction.

15 **[00221]** If force option is chosen, the operation can succeed if not all CDs are available, but a warning to the user as well as a log entry identifying the set of "missing" CDs can be made if not all CDs are available.

20 **[00222]** In order for the *start* command to succeed, the enabled status for the use of secret M_Key may need to be set. The enabled status can initially be set via using the *enablesecretmkey* command, or via choosing the "enable" option that can be specified for the *start* command. The reason for having the "enable" option for the *start* command is to ensure that secret M_Key usage can be enabled together with a defined current value and/or list of known values.)

25 **[00223]** Furthermore, if "addonly" type is specified for a transaction, then no delete operations is allowed in this transaction (vice versa if "deleteonly" is specified). It is also not legal to specify both "deleteonly" and "addonly" for the same transaction. If neither "deleteonly" nor "addonly" is specified, then only a *set-current* operation is allowed. On the other hand, the *set-current* operation can be used independently of whether the transaction is "addonly" or "deleteonly".

listactive

30 **[00224]** The *listactive* command can display the last committed list of M_Keys as well as the current M_Key value and the enabled status. This command can be invoked independently of any on-going transaction. The *listactive* command can include options for listing symbolic and/or hex.

listworking

35 **[00225]** The *listworking* command can display the resulting list and current base of changes since the last *start* command (i.e. that has not been committed). The *listworking* command can

include options for listing symbolic and/or hex.

listlocalmkey

5 **[00226]** The *listlocalmkey* command can display the current M_Key value (if any) defined for local switch chip (e.g. an IS4 instance). The output can reflect if any M_Key is defined, and if so if it is "readable" or "secret". Additionally, the output can reflect the case where no M_Key is used.

[00227] In the case of configuration errors (e.g. provoked by manual access to the sysfs interface), the output can also reflect illegal, or un-supported, protection bit values. Also, the output can reflect that the NM2 SM is not able to manage this node, e.g. when the M_Key value
10 is zero, but the protection bit value indicates a "secret" mode.

[00228] This operation can be performed from any NM2 independently of local SM enabled or master status.

setlocalsecretmkey

15 **[00229]** The *setlocalsecretmkey* command can set the M_Key value for the local IS4 instance to the specified value and set the associated mode to "secret" (i.e. the M_Key protect bits). The *setlocalsecretmkey* command can include options for specifying symbolic with password requirements or plain hex.

[00230] This operation can be performed from any NM2 node, independently of local SM
20 enabled or master status. The value zero is not a legal secret M_Key value, and the CLI command can reject it.

add

[00231] The *add* command can add a new M_Key value to the working list. The *add*
25 command can include options for specifying symbolic with password requirements or plain hex. The *add* command is only legal if the transaction has been started with "addonly" option.

[00232] The revoking of any incorrectly added value(s) can be accomplished by explicitly aborting the current transaction via the abort subcommand. This ensures that the adding and
then
30 the setting current of a newly added value is cleaned up in a consistent and well defined way.

delete

[00233] The *delete* command can remove an existing M_Key value from the list. The *delete*
command can include options for specifying symbolic or plain hex.

35 **[00234]** This command is only legal if the transaction is started with "deleteonly" option. Furthermore, it is not legal to remove the "current" value, nor the "old current" value, which is

enforced by the interface. Also, it is not legal to delete any value that may still exist as "current" for any NM2 node in the system that may be temporarily offline at the moment. This constraint may not be enforced by the interface, but can be documented as a rule that the administrator should follow.

- 5 **[00235]** Additionally, the revoking of any incorrectly deleted value(s) can be accomplished by explicitly aborting the current transaction via the *abort* subcommand.

set-current

- 10 **[00236]** The *set-current* command can set an already defined M_Key value as the current M_Key value that is used by the master SM following a subsequent commit operation. The *set-current* command can include options for specifying symbolic or plain hex.

- 15 **[00237]** When updating current M_Key value, then the old current M_Key value can be recorded. It is not legal to both update the current M_Key value and delete the old current M_Key value in the same update transaction. Also, multiple *set-current* operations during the same transaction can be legal as long as they all identify a currently defined M_Key value, and only the last operation prior to commit may have any impact on the resulting active configuration.

[00238] If no current M_Key value is defined when a configuration is committed, then this is accepted, but the result is the same as if the configuration had been disabled.

- 20 **[00239]** Once a current M_Key value is defined via the *set-current* command, then it is not possible to unset it unless the configuration is disabled. This also implies that it is not possible to add more known values while a configuration is disabled.

[00240] Thus, in order to update the known list when in disabled mode, the configuration has to be temporarily enabled while the configuration is updated, but can then be immediately disabled again afterwards.

25

commit

[00241] The *commit* command can make the configuration that has been (re)defined since the last *start* operation active in the system via the PD-PD commit protocol. The *commit* command can include option for force commit.

- 30 **[00242]** By default, the operation can fail if not all defined ("smnodes") PD instances are present. If "force" option is specified, then the operation may still succeed if not all PD instances are available but a warning and log message can be generated.

abort

- 35 **[00243]** The *abort* command can discard all operations that have not been committed since the last *start* command without any impact on the active configuration. Any incorrect argument to *delete* or *add* commands can be revoked by performing an *abort* operation.

setpassword

[00244] The *setpassword* command can define a password to be used for encrypting M_Key list before replicating to other nodes. Furthermore, same password can be defined on all "smnodes" NM2s and can be updated in the same controlled way as when updating "smnodes" membership.

[00245] The password with an alphanumeric string of at least 8 characters is considered strong. The system may provide a built-in-default password value, which can be replaced via the *setpassword* command.

[00246] The method for encryption can be based on an open source toolkit, where the user-supplied password can be used as input to the encryption/decryption algorithms.

[00247] The specified password can also be used to generate derived key values based on key values specified by the *add* command. The specified password can be persistently recorded locally on the NM2, until it is explicitly updated again.

[00248] Whenever a password has been updated on one "smnode", it has to be updated on the other "smnodes" before any transaction can take place.

enablesecretmkey

[00249] The *enablesecretmkey* command can represent an implicit start and commit transaction for the currently defined configuration. The *enablesecretmkey* command can include an option for "force start/commit".

[00250] If no configuration is currently defined, then the command may have no effect, except for verifying that all "smnodes" peers are available and in the correct state, and also setting the enabled status to "enabled". (The enabled status is a dedicated flag in the underlying configuration file.) If no configuration is currently defined, then the enabled status may have no effect on the operation of the SM.

disablesecretmkey

[00251] The *disablesecretmkey* command can represent an implicit start /commit transaction for the currently defined configuration with the side effect that no secret M_Key value can be used by any SM on "smnodes". Unlike an ordinary transaction, the start /commit transaction does not update the current M_Key value defined by the configuration, but sets the "enabled status" to "disabled". The *disablesecretmkey* command can include an option for "force start/commit".

[00252] The disabled status implies that the SM may either be using any already defined readable M_Key value, or not use any M_Key value at all. Also, the list of known secret M_Key values can still be defined, and the master SM can still be able to discover and manage any

port/node that was not available during the initial update to "non-secret" M_Key, but that is now available and with one of the known secret M_Key values active.

[00253] The update transactions, e.g. the transactions caused by *commit*, *enablesecretmkey* or *disablesecretmkey* subcommands, can be implemented in a way that is safe in terms of that
5 the master election state within an operational subnet with NM2 2.1 firmware on all SM enabled NM2s can be well defined also if the transaction fails.

[00254] Furthermore, a failed transaction may be repeated in order to try to complete it successfully in the case of a transient problem without any inherent risk of more dramatic inconsistencies. Also, fencing operations, such as defining an additional known key value as well
10 as making it current, can be done in force mode to ensure that an NM2 node with unknown state may not be able to modify the state of the subnet.

[00255] Additionally, the disabling use of secret keys following a failed update transaction, or the deletion of known key values following a failed transaction may lead to inconsistent master election in the subnet and may therefore not be attempted before the set of smnodes have not
15 been brought to a consistent state.

[00256] Here, bringing the system to a consistent state can involve performing a *enablesecretmkey* or an update transaction from the current master, based on whether the (possibly new) master has the updated (desired) current configuration.

Claims:

What is claimed is:

- 5 1. A method for supporting network management in a network environment, comprising:
receiving, via a configuration daemon (CD) on a switch in the network environment, a
configuration file, wherein the configuration file includes a list of known management key
(M_Key) values;
storing the configuration file, and making the configuration file available to a local subnet
10 manager(SM) on the switch, wherein the local SM is associated with an currently used M_Key
value; and
updating the local SM with a new M_Key, after receiving an instruction from a master CD
that is associated with a master SM in the network environment.
- 15 2. The method according to Claim 1, further comprising:
including in the configuration file a trusted M_Key and one or more key values that are
transparently generated based on the trusted M_Key.
3. The method according to Claims 1 or 2, further comprising:
20 continuing using the old current M_Key when updating/setting M_Key value for any
subnet management agent (SMA) instance until the local SM is updated with the new M_Key.
4. The method according to any preceding Claim, further comprising:
using, via the CD, a dedicated SM command line interface (CLI) command to update the
25 local SM.
5. The method according to any preceding Claim, further comprising:
updating, via the CD, the local SM with current local secret key configuration when at
least one of the SM and CD is restarted.
- 30 6. The method according to any preceding Claim, further comprising:
containing in the configuration file information about the currently used M_Key value as
well as the intended new M_Key value.
- 35 7. The method according to any preceding Claim, further comprising:
recording the currently used M_Key value in the list of known M_Key values.

8. The method according to any preceding Claim, further comprising:
using the new M_Key value to fence off a SM on an unavailable node.
9. The method according to any preceding Claim, further comprising:
5 ensuring that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.
10. The method according to any preceding Claim, further comprising:
initiating a commit operation on the master CD.
- 10
11. A computer program comprising program instructions for running on one or more microprocessors to perform all the steps of the method of any preceding claim.
12. A computer program comprising the computer program of claim 11 provided on a
15 machine-readable medium.
13. A system for supporting network management in a network environment, comprising:
one or more microprocessors;
a configuration daemon (CD) on a switch running on the one or more microprocessors,
20 wherein the CD operates to
receive a configuration file that includes a list of known management key (M_Key) values;
store the configuration file, and make the configuration file available to a local subnet manager(SM) on the switch, wherein the local SM is associated with an currently
25 used M_Key value; and
update the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.
14. The system according to Claim 13, wherein:
30 the configuration file includes a trusted M_Key and one or more key values that are transparently generated based on the trusted M_Key.
15. The system according to Claims 13 or 14, wherein:
the local SM continues using the old current M_Key when updating/setting M_Key value
35 for any subnet management agent (SMA) instance until the local SM is updated with the new M_Key.

16. The system according to any of Claims 13-15, wherein:
the CD operates to use a dedicated SM command line interface (CLI) command to update the local SM.
- 5 17. The system according to any of Claims 13-16, wherein:
the CD operates to update the local SM with current local secret key configuration when at least one of the SM and CD is restarted.
18. The system according to any of Claims 13-17, wherein:
10 the configuration file contains information about the currently used M_Key value as well as the intended new M_Key value.
19. The system according to any of Claims 13-18, wherein:
the CD operates to record the currently used M_Key value in the list of known M_Key
15 values.
20. The system according to any of Claims 13-19, wherein:
the new M_Key value is used to fence off a SM on an unavailable node.
- 20 21. The system according to any of Claims 13-20, wherein:
the master CD operates to ensure that the master SM has a known M_Key list not shorter than any M_Key list maintained by a standby SM instance.
22. A non-transitory machine readable storage medium having instructions stored thereon
25 that when executed cause a system to perform the steps comprising:
receiving, via a configuration daemon (CD) on a switch in a network environment, a configuration file, wherein the configuration file includes a list of known management key (M_Key) values;
storing the configuration file, and making the configuration file available to a local subnet
30 manager(SM) on the switch, wherein the local SM is associated with an currently used M_Key value; and
updating the local SM with a new M_Key, after receiving an instruction from a master CD that is associated with a master SM in the network environment.
- 35 23. A computer program for causing a computer to implement the method recited in any one of Claims 1 to 10.

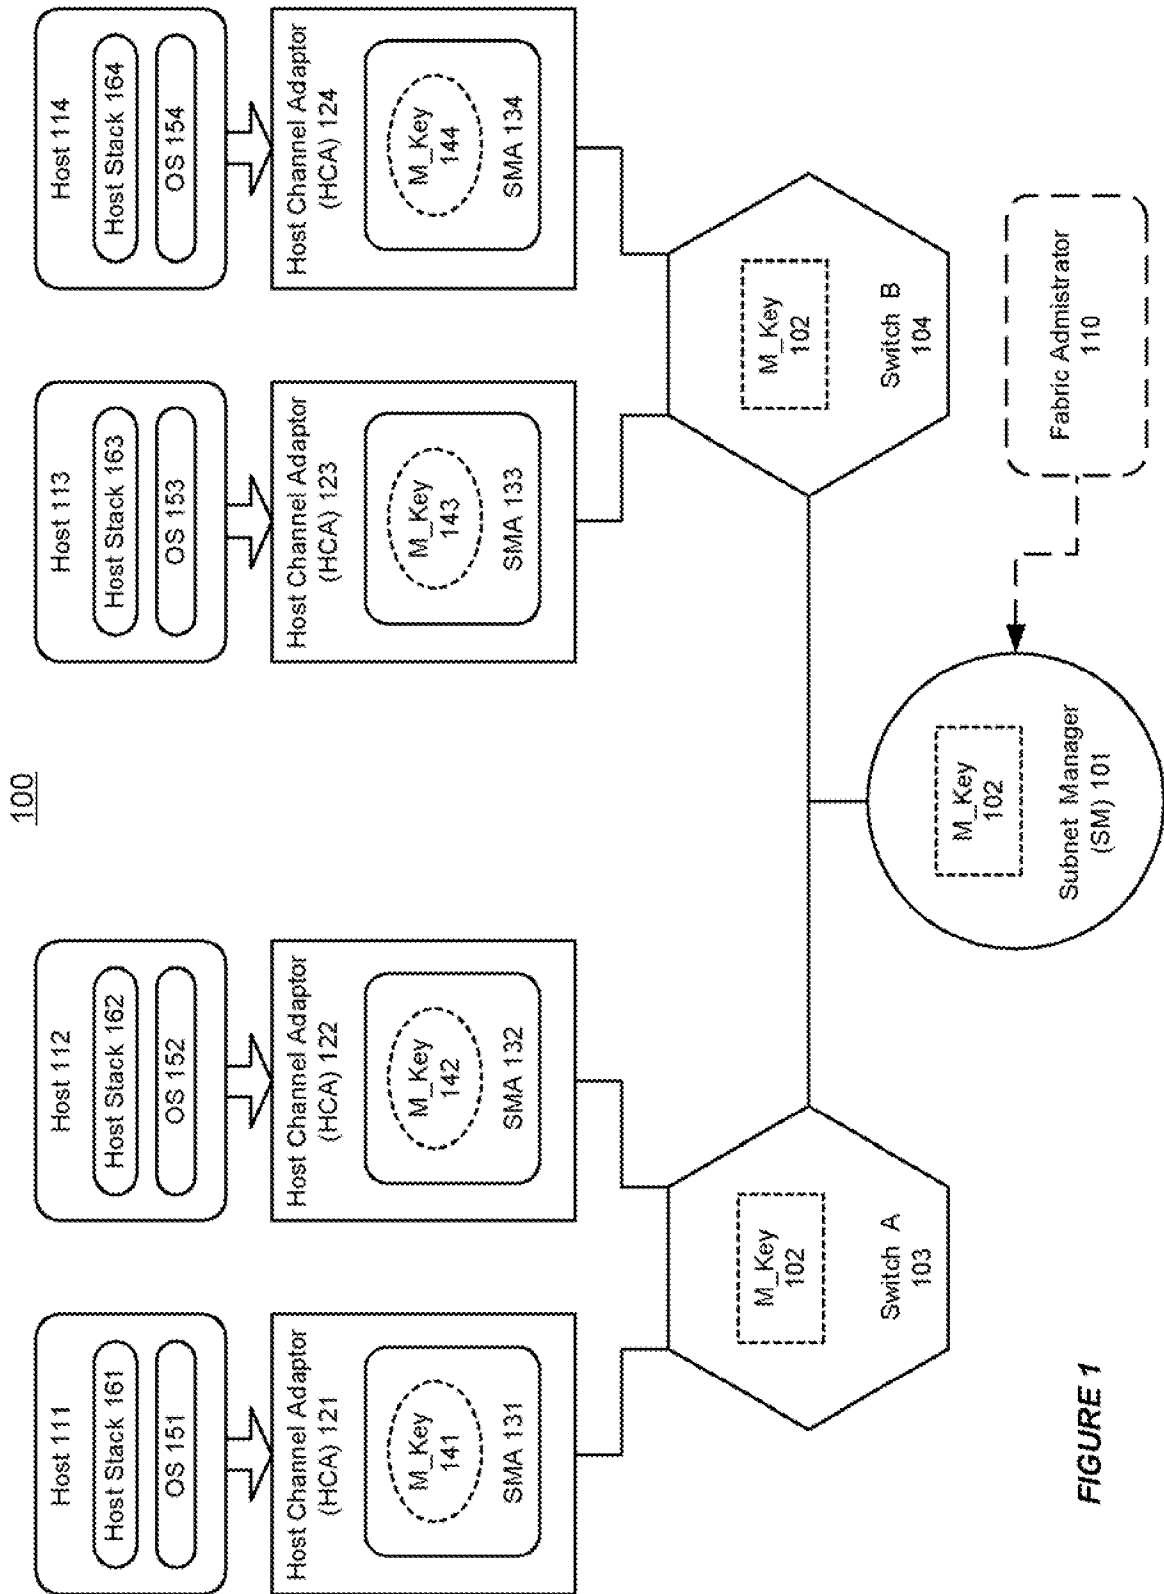


FIGURE 1

200

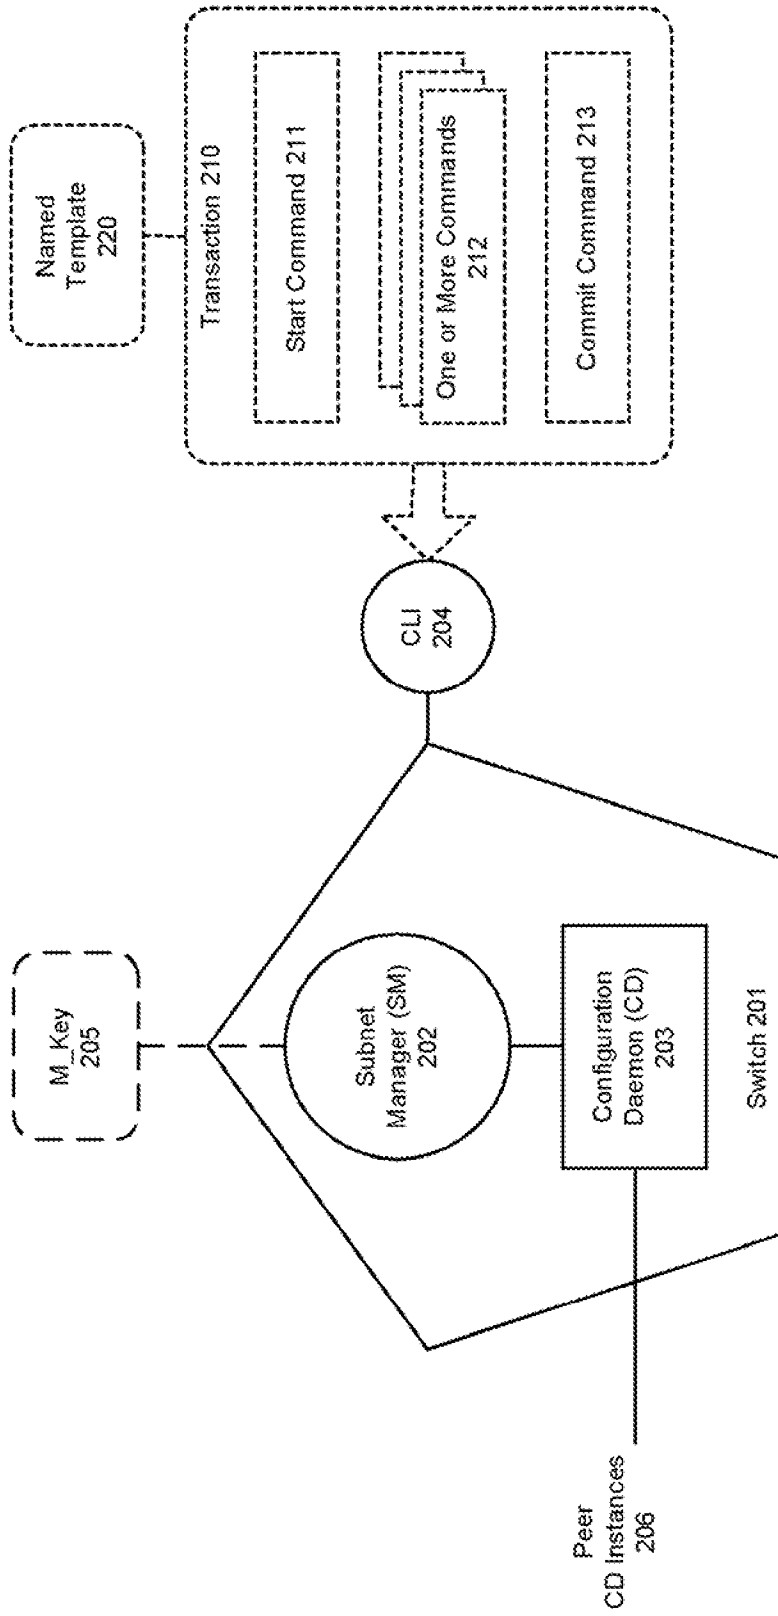


FIGURE 2

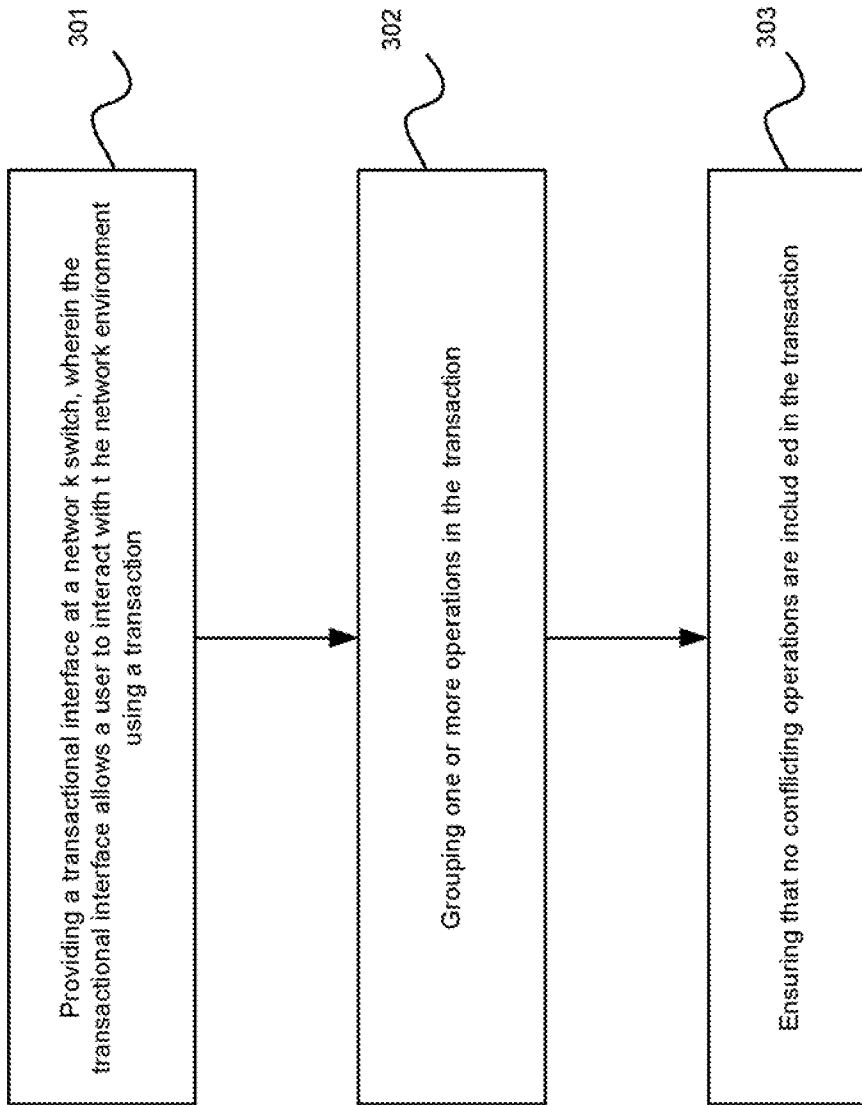


FIGURE 3

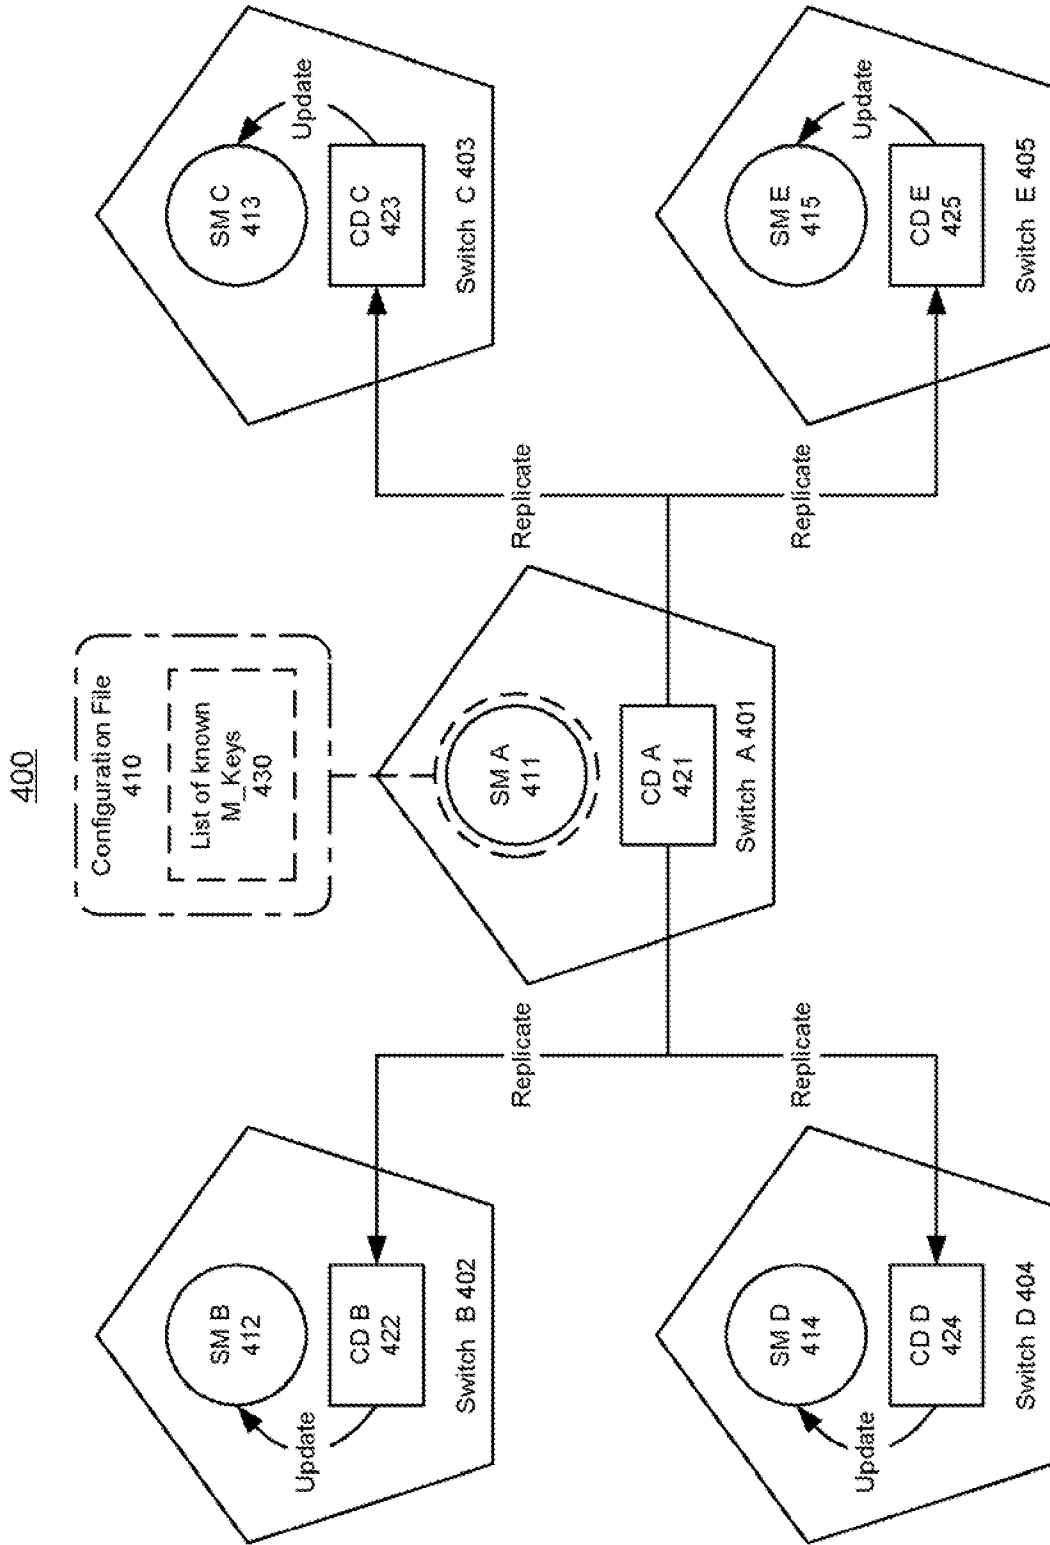


FIGURE 4

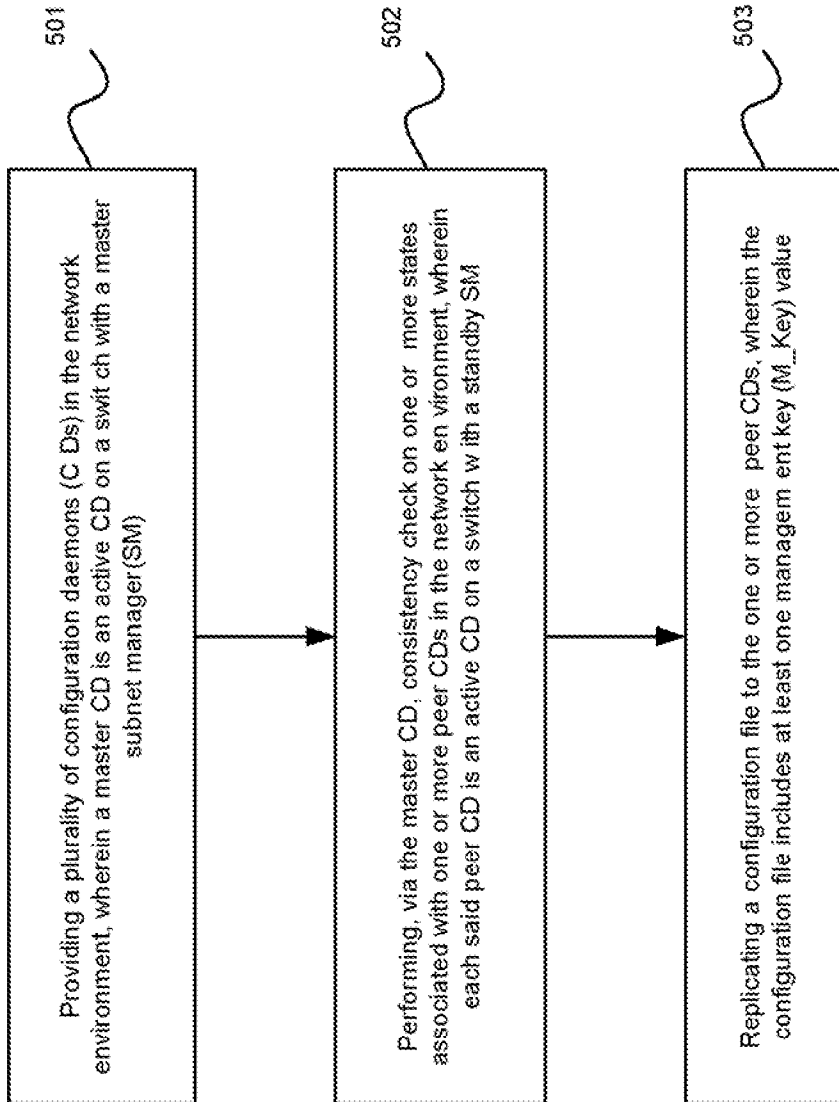


FIGURE 5

600




FIGURE 6




FIGURE 7

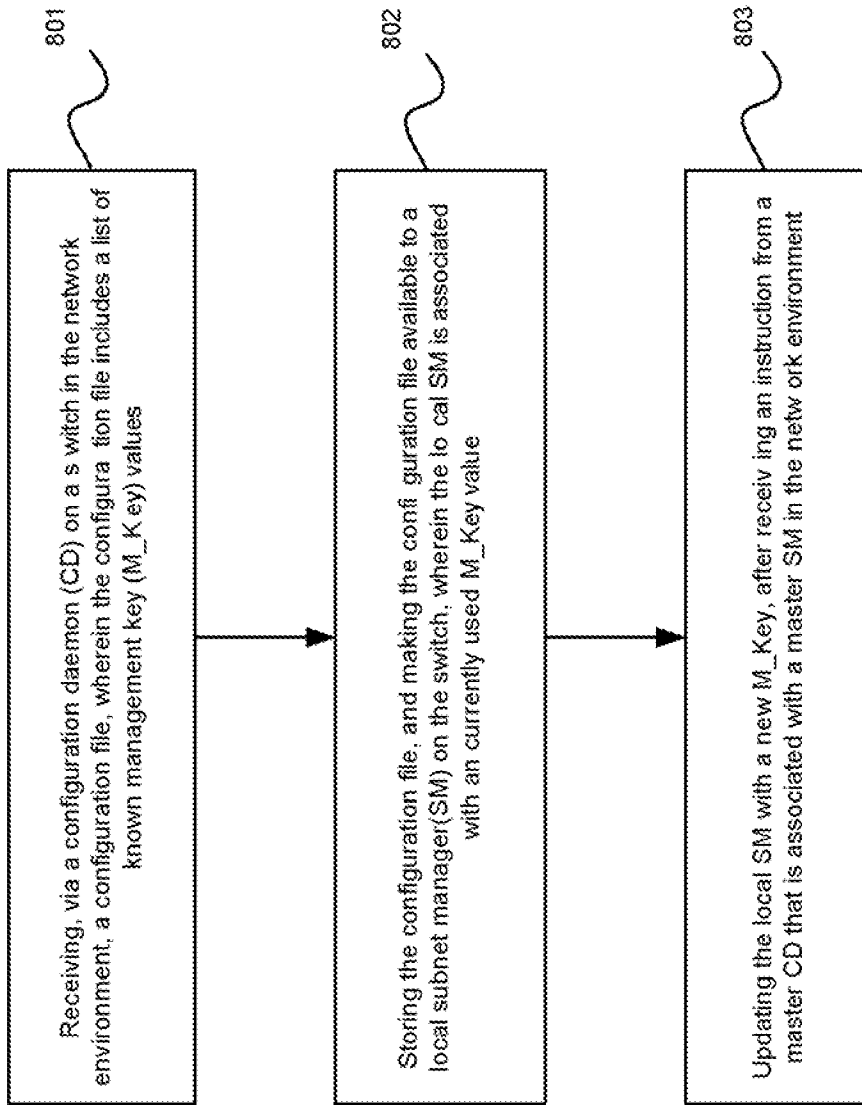


FIGURE 8




FIGURE 9

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/040639

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H04L29/06 H04L29/08
 ADD.
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 H04L

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 EPO-Internal, WPI Data, COMPENDEX, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 7 290 277 B1 (CHOU NORMAN C [US] ET AL) 30 October 2007 (2007-10-30) column 3, line 49 - column 6, line 42 column 7, line 1 - column 8, line 11 column 9, line 53 - column 10, line 25 figures 1-5	1-23
A	----- WO 2012/037518 A1 (ORACLE INT CORP [US]; JOHNSEN BJORN-DAG [NO]; HOLEN LINE [NO]; MOXNES) 22 March 2012 (2012-03-22) paragraphs [0055] - [0066] figures 1-10	1-23
A	----- US 6 941 350 B1 (FRAZIER GILES ROGER [US] ET AL) 6 September 2005 (2005-09-06) column 9, line 38 - column 11, line 20 figures 6-8	1-23

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search 13 September 2013	Date of mailing of the international search report 23/09/2013
---	---

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Ghomrasseni, Z
--	---

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2013/040639

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7290277	B1	30-10-2007	NONE

WO 2012037518	A1	22-03-2012	CN 103125097 A 29-05-2013
			CN 103125098 A 29-05-2013
			CN 103125102 A 29-05-2013
			EP 2617157 A1 24-07-2013
			EP 2617159 A1 24-07-2013
			EP 2617165 A1 24-07-2013
			US 2012069730 A1 22-03-2012
			US 2012072562 A1 22-03-2012
			US 2012072563 A1 22-03-2012
			US 2012072564 A1 22-03-2012
			US 2012079090 A1 29-03-2012
			US 2012079580 A1 29-03-2012
			WO 2012037512 A1 22-03-2012
			WO 2012037518 A1 22-03-2012
			WO 2012037520 A1 22-03-2012

US 6941350	B1	06-09-2005	NONE
