

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
2 April 2009 (02.04.2009)

PCT

(10) International Publication Number
WO 2009/042856 A1

- (51) International Patent Classification:
H04L 12/403 (2006.01)
- (21) International Application Number:
PCT/US2008/077838
- (22) International Filing Date:
26 September 2008 (26.09.2008)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/960,444 28 September 2007 (28.09.2007) US
- (71) Applicant (for all designated States except US): **ALLIED
TELESIS, INC.** [US/US]; 19800 North Creek Parkway,
Suite 200, Bothell, WA 98011 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **FOWLER, Hayden**
[NZ/NZ]; c/o Allied Telesis, Inc., 19800 North Creek Park-
way, Suite 200, Bothell, WA 98011 (US). **BEALE, Tim**
[NZ/NZ]; c/o Allied Telesis, Inc., 19800 North Creek Park-
way, Suite 200, Bothell, WA 98011 (US). **MCNAB, Sha-
ree** [NZ/NZ]; c/o Allied Telesis, Inc., 19800 North Creek
Parkway, Suite 200, Bothell, WA 98011 (US). **SATOH,**

Seiichiro [JP/JP]; c/o Allied Telesis, Inc., 19800 North
Creek Parkway, Suite 200, Bothell, WA 98011 (US).

- (74) Agents: **HAYDOUTOVA, Juliana** et al.; Arent Fox LLP,
1050 Connecticut Avenue, N.W., Suite 400, Washington,
DC 20036 (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, BR, BW, BY, BZ, CA,
CH, 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, PG, PH, PL, PT,
RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL,
NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG,
CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR PREVENTING NETWORK CONFLICT




Figure 2A

(57) Abstract: A resiliency-link having a master node and slave nodes may be used to automatically avoid network conflicts due to failure of a node stack. Any stack stub formed as a result of the failure of a stack node or link may use the resiliency-link to determine whether the stack Master node is still online and providing service. If communication over the resiliency-link determines that the original Master node is offline, a new stack stub takes over and provides the desired service. If the original Master node is still online, then the new stack stub takes an appropriate alternative or fallback action.

WO 2009/042856 A1



Published:

— *with international search report*

METHOD AND APPARATUS FOR PREVENTING NETWORK CONFLICT

BACKGROUND OF THE INVENTION

[1] This application claims priority from provisional U.S. Patent Application No. 60/960,444, filed on September 28, 2007, titled "Virtual Chassis Stacking." This provisional application is incorporated herein in its entirety.

1. Field of Invention

[2] The present invention relates to an apparatus and method to prevent network conflicts. More specifically, the present invention relates to an apparatus and method to prevent network conflicts in case of link or node failure in a network stack.

2. Description of Related Art

[3] A network comprises a number of network nodes. A node stack provides a means to group one or more of these network nodes into a single virtual network node that may provide a desired service, for example hosting an IP address. When a failure occurs, which generally results in the stack nodes becoming disconnected from each other, stack stubs are created, which are smaller groups of nodes that are still connected to each other after the failure has occurred. Thus, a newly formed stack stub may no longer be able to use the stack's preferred communication path to determine if the Master node (e.g., in another stub stack) is still online and providing service, or whether the Master node is no longer operational.

[4] As a result, an undesirable outcome of the failure may be that any resulting stack stubs may end up providing duplicate services with the original stack. For example, a first stub and a distinct second stub may both end up hosting the same IP address that was hosted by the original stack. A network conflict arises as a result. Alternatively, an equally undesirable outcome occurs when none of the stack stubs resulting from the failure provides the service, which results in a failure of the service. In the above example, the result would be that the IP address becomes unreachable.

[5] There is a need in the art, therefore, for an apparatus and method for preventing network failures that result in duplication of services or failures to provide services.

SUMMARY OF THE INVENTION

[6] In light of the above described problems and unmet needs and according to various exemplary aspects of the current invention, a resiliency-link apparatus, and method of operation thereof, provides an alternative communication path used for monitoring the health of a given stack and the nodes that form the stack. The health of a stack node may be, for example, whether the node is online or offline, or whether the stack is capable of providing a particular service such as hosting an IP address. The various exemplary aspects of the current invention may facilitate dynamic and automated recovery from catastrophic stack failures, where the preferred communication path(s) between stack nodes is severed.

[7] According to various exemplary aspects of the current invention, a resiliency-link may be used to overcome the above-described undesirable situations by providing a stacking solution that is highly resilient to failures. The resiliency link may be similar to a network link, but with a bandwidth that is lower than a typical network link, and may contain only stack control information. Any new stack stub, formed as a result of the failure of a stack node or of a link, may use the resiliency-link to determine whether the original stack stub is still online and providing service. According to various exemplary aspects of the current invention, if, via a communication over the resiliency-link, it is determined that the original stack stub is offline, the new stack stub takes over and provides the desired service such as, for example, hosting an IP address. On the other hand, if via a communication over the resiliency-link, it is determined that the original stack master is online, then the new stack master may take an appropriate alternative or fallback action. According to various exemplary aspects of the current invention, that appropriate or alternative fallback action may be, for example, hosting a different IP address, or not hosting any IP address at all. It should be noted that the service provided by a stack is defined by a configuration. Thus, configuring a stack may be providing an IP address, for example, which is the service provided by the stack.

[8] According to various exemplary aspects of the current invention, the resiliency-link may be used in conjunction with a fallback-configuration, which may dynamically reconfigure a new stack stub resulting from a failure within an original stack, to provide alternative services that do not conflict with the services provided by the original stack. The resiliency-link and fallback-configuration may allow a stack to recover from failures, without user intervention, by allowing the stack to split into two or more separate stack stubs, which may co-exist in a stable condition, thereby avoiding network conflict.

[9] According to various exemplary aspects of the current invention, a resiliency link may be used with a ring topology stack, a star topology, a dual-star topology, a mesh topology, or any other stack topology. According to various exemplary aspects of the current invention, the resiliency link may be a secondary ring connecting the various nodes within a given stack, or may use a hub to connect to each node of a given stack.

[10] Additional advantages and novel features of these aspects of the invention will be set forth in part in the description that follows, and in part will become more apparent to those skilled in the art upon examination of the following or upon learning by practice of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[11] Various exemplary embodiments of the systems and methods will be described in detail, with reference to the following figures, wherein:

[12] Figs. 1A-1D are diagrams illustrating conventional node stacks experiencing failure;

[13] Figs. 2A-2E are diagrams illustrating node stacks including resiliency links, according to various aspects of the current invention;

[14] Fig. 3 is a flow chart illustrating a method of preventing network conflict in a node stack, according to various aspects of the current invention;

[15] Fig. 4 presents an exemplary system diagram of various hardware components and other features, for use in accordance with aspects of the present invention; and

[16] FIG. 5 is a block diagram of various exemplary system components, in accordance with aspects of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[17] These and other features and advantages of this invention are described in, or are apparent from, the following detailed description of various exemplary aspects of the current invention.

[18] Figs. 1A-1D are diagrams illustrating conventional node stacks experiencing failure. Fig. 1A illustrates a conventional node stack that includes a Master node A and several slave nodes B, C, D and E in a ring topology stack, in which a failure of two data-links has occurred. In Fig. 1A, the data or network links between nodes B and C, and the data or network links between nodes D and E have failed and can no longer transmit data. As a result, two stack stubs are formed, a first stub including the original Master node A with nodes B and E, and a second stack stub is formed that includes nodes C and D. In the second stack stub, upon detection of the failure of the original stack, a new stack master is elected, where one of the nodes C or D becomes a master node and the new master node C or D will have the same configuration as the original Master node A. In other words, the new master node C or D will configure itself to provide the same services as Master node A. As a result, two stack stubs are created, each having the same configuration, which results in a network conflict.

[19] Fig. 1B illustrates a conventional node stack including a Master node A and several slave nodes B, C, D and E in a ring topology stack, in which the data link between nodes D and E has failed, and node B has failed. As a result, two stack stubs are formed, a first stub including the original Master node A and node E, and a second stack stub is formed that includes nodes C and D. In the second stack stub, upon detection of the failure in the original stack, a new stack master is elected, where one of the nodes C or D becomes a master node, and the new master node C

or D will have the same configuration as the original Master node A. In other words, the new master node C or D will configure itself to provide the same services as Master node A. Accordingly, two stack stubs are created, both having the same configuration, which results in a network conflict.

[20] Fig. 1C illustrates a conventional node stack including a Master node A and several slave nodes B, C, D and E in a star topology stack, in which the data link between nodes A and D has failed. As a result, two stack stubs are formed, a first stub that includes the original Master node A and nodes B, C and E, and a second stack stub that is formed by node D. In the second stack stub formed by node D, upon detection of the failure in the original stack, D elects itself as a new stack master, and the new master node D will have the same configuration, and/or provide the same service, as the original Master node A. Accordingly, two stack stubs are created, both having the same configuration, which results in a network conflict.

[21] Fig. 1D illustrates a conventional node stack including a Master node A and several slave nodes B, C, D, E and F in a dual-star topology stack, in which the data links between nodes A and D, and between nodes F and D, have failed. As a result, two stack stubs are formed, a first stub that includes the original Master node A and nodes B, C, E and F, and a second stack stub that is formed by node D. In the second stack stub formed by the node D, upon detection of the failure in the original stack, D elects itself as a new stack master, and the new master node D will have the same configuration, and/or provide the same service, as the original Master node A. Accordingly, two stack stubs are created, both having the same configuration, which results in a network conflict.

[22] Figs. 2A-2E are diagrams illustrating node stacks including resiliency links, according to various aspects of the current invention. Fig. 2A is a diagram illustrating a node stack including a Master node A and several slave nodes B, C, D and E in a ring topology stack according to various exemplary aspects of the present invention, in which the data links between nodes B and C, and between nodes D and E have failed. According to various exemplary aspects of the current invention, the node stack also includes a hub H that is connected to each of nodes A, B, C, D and E, via a resiliency link. The resiliency link may be a network link having bandwidth

that is lower than the bandwidth of a typical network data link. The resiliency link may contain stack control information only, and the resiliency link may be devoid of user data that is typically carried on a network data link. The resiliency link may also be less expensive than a conventional network data link forming a ring. Upon failure of the data links between nodes D and E, and between nodes B and C, two stack stubs are formed. The Master node A forms a first stack stub with nodes B and E, and nodes C and D form a second stack stub. Nodes C and D would negotiate to select a master node for the second stack stub. In the absence of the resiliency link, nodes C and D would apply the original stack configuration. However, because the resiliency link connects each of the nodes A, B, C, D and E to the hub H, the hub H may convey the information to each of the nodes B, C, D and E that the Master node A is still operational. When nodes C and D in the second stub detect, via the resiliency link, that the original Master node A is still operational, the nodes C and D may load an alternative configuration from an alternative configuration file to prevent a network configuration conflict. Alternatively, nodes C and D may load no configuration at all. In either case, no network conflict results.

[23] Fig. 2B is a diagram illustrating a node stack including a Master node A and several slave nodes B, C, D and E in a ring topology stack without a hub, according to various exemplary aspects of the present invention. In Fig. 2A, the data links between nodes B and C and between nodes D and E have failed. According to various exemplary aspects of the current invention, the node stack also includes a resiliency link R that connects nodes A, B, C, D and E. The resiliency link R may be a network link having a bandwidth that is lower than the bandwidth of a typical network data link, and may only contain stack control information. The resiliency link R may be devoid of user data that is typically transmitted over a network data link. Upon failure of the data links between nodes D and E, and between nodes B and C, two stack stubs are formed, but the resiliency link R remains connected to all the nodes A, B, C, D and E. The Master node A forms a first stack stub with nodes B and E, and nodes C and D form a second stack stub. Nodes C and D would negotiate to select a master node for the second stack stub. In the absence of the resiliency link R, nodes C and D would apply the original stack configuration. However, because the resiliency link R connects each of nodes A, B, C, D and E together, the resiliency link R may convey the information to each of the

slave nodes, including nodes C and D, that Master node A is still operational. When nodes C and D in the second stack stub detect, via the resiliency link R, that the original master node A is still operating, nodes C and D may load an alternative configuration to prevent a network configuration conflict. Alternatively, nodes C and D may load no configuration at all. In either case, no network conflict results.

[24] Fig. 2C is a diagram illustrating a node stack including a Master node A and several slave nodes B, C, D and E in a star topology stack according to various exemplary aspects of the present invention. In Fig. 2C, the data link between nodes A and B has failed. According to various exemplary aspects of the current invention, the node stack also includes a hub H that is connected to each of the nodes A, B, C, D and E, via a resiliency link. The resiliency link may be a network link having a bandwidth that is lower than the bandwidth of a typical network data link, and may only contain stack control information. The resiliency link may be devoid of user data that is typically transmitted over a network data link. Upon failure of the data link between nodes A and B, two stack stubs are formed. Master node A forms a first stack stub with nodes C, D and E, and node B alone forms a second stub stack. Node B would become the master node for the second stack stub. In the absence of the resiliency link, node B would apply the original stack configuration. However, because the resiliency link connects node B to the hub H, the hub H may convey the information to node B that the Master node A is still operational. When node B in the second stub detects, via the resiliency link, that the original master node A is still operating, node B may load an alternative configuration to prevent a network configuration conflict. Alternatively, node B may load no configuration at all. In either case, no network conflict results.

[25] Fig. 2D is a diagram illustrating a node stack including a Master node A and several slave nodes B, C, D, E and F in a dual-star topology stack, in which the data links between nodes A and B and between nodes B and F have failed. According to various exemplary aspects of the current invention, the node stack also includes a hub H that is connected to each of the nodes A, B, C, D, E and F via a resiliency link. The resiliency link may be a network link having a bandwidth that is lower than the bandwidth of a typical network data link, and may only contain stack control information. The resiliency link may be devoid of user data that is typically

transmitted over a network data link. Upon failure of the data links between nodes A and B and between nodes B and F, two stack stubs are formed. The Master node A forms a first stack stub with nodes C, D, E and F, and node B alone forms a second stack stub. Node B would negotiate to select a master node for the second stack stub. In the absence of the resiliency link, node B would apply the original stack configuration. However, because the resiliency link connects node B to the hub H, hub H may convey the information to node B via the resiliency link that the Master node A is still operational. When node B detects, via the resiliency link, that the original master node A is still operational, node B may load an alternative configuration to prevent a network configuration conflict. Alternatively, node B may load no configuration at all. In either case, no network conflict results.

[26] Fig. 2E is a diagram illustrating a node stack including a Master node A and several slave nodes B, C, D and E in a ring topology stack, in which the data link between nodes C and D, and node A, have failed. According to various exemplary aspects of the current invention, the node stack also includes a hub H, connected to each of nodes A, B, C, D and E, via a resiliency link. The resiliency link may be a network link having a bandwidth that is lower than the bandwidth of a typical network data link, and may only contain stack control information. The resiliency link may be devoid of user data that is typically transmitted over a network data link. Upon failure of the data link between nodes C and D, and failure of node A, two stack stubs are formed, a first stack stub is formed with nodes D and E, and a second stack stub is formed with nodes B and C. In this case, and in the absence of a master node, nodes B and C negotiate to elect a master node for the first stub, and nodes D and E negotiate to elect a master node for the second stub. Hub H conveys information regarding the operative state of the nodes to each of the nodes via the resiliency links. The resiliency link here may be used to allow negotiation to take place between the two stack stubs to determine which stack stub will apply the preferred configuration and which stack stub will apply an alternate configuration, by conveying information via the hub H. Accordingly, no network conflict results.

[27] According to various aspects of the current invention, node failure may be detected in various ways. One of the ways in which failure is detected is through the use of a periodic multicast message sent to the nodes of the node stack. For

example, a multicast message may be sent regularly, e.g., every 500 ms, over the resiliency link from the master node to the slave nodes. According to various aspects of the current invention, the message may contain information such as the master node's MAC address and an incrementing sequence number generated by the master node. The broadcast messages can be construed to constitute health-check messages. As long as a slave node receives these health-check messages from the master node, the slave node considers the master node to be online and operational. If no health-check messages are received from the master node by a slave node over a predetermined period of time, for example, a period of two seconds, then the slave node considers the master node to be offline and inoperative. According to various exemplary aspects of the current invention, when the master node becomes unreachable over the network data links, the one or more slave nodes that can no longer reach the master node form a stack stub, and elect a new master. If the new master no longer receives the health-check messages, indicating that the original master node is no longer operational, then the new master node may load the original configuration. On the other hand, if the health-check messages do not time out, then the new master may reconfigure the stub stack using a fallback and/or alternative configuration. According to various aspects of the current invention, a fallback configuration may range from specifying a back-up IP address to a complete configuration for an alternative stack setup. If no fallback configuration is available, then the slave nodes may deconfigure themselves and shut down all network interfaces to avoid a network conflict.

[28] It should be noted that other communication protocols can be used, according to aspects of the current invention, over the resiliency links to ascertain the status of a stack node. Such communication protocols may be, for example, Internet Control Message Protocol (ICMP) echo requests, Address Resolution Protocol (ARP)/ Reverse Address Resolution Protocol (RARP) requests, or discovery protocols such as Link Layer Discovery Protocol (LLDP).

[29] The health-check messages may also, according to various aspects of the current invention, specify a dynamic negotiation of the recovery action for a given stack. For example, the fallback configuration may be negotiated based upon which stack stub has the least number of nodes present, so that the preferred stack stub

configuration is applied to the stack stub that has the most bandwidth. Alternatively, the stack stubs may have several fallback configuration files available, and negotiate which configuration file should be applied to the new stack.

[30] Fig. 3 is a flow chart illustrating a method of preventing network conflict in a node stack, according to various aspects of the current invention. The method starts at S110, where a node stack operating to provide a given service such as, for example, hosting an IP address, detects a link or node failure. According to various exemplary aspects of the current invention, the node stack includes a resiliency link, which may be similar to a secondary ring, and may have a bandwidth that is smaller than or equal to the bandwidth of a conventional data ring. The resiliency link connects each node in the node stack to at least one other node in the node stack. Alternatively, the resiliency link may be connected to a hub that connects each node of the node stack. Next, the method continues to S120, where a check is made to see if the stack master is still connected to each slave node over the network data links. In S130, in the event that the Master is still reachable from every slave node over the network data links, no isolated stacking stub has been formed, and the method continues to S140, where normal stack operation continues and no stub master is elected.

[31] In S130, in the event that the master is not reachable from each slave node, the method continues to S150 to elect a new stack master for an isolated stacking stub that has been formed as a result. The method then continues to S160, where a check may be made to determine whether the Master health check message has been received by every slave node over the resiliency link. If the Master node has failed and no health check messages have been received, then the method continues to S170, where the original configuration is applied to the local stub master. If in S160 the Master node has not failed, then the method continues to S180, where a backup configuration is applied to the local stub master, and there may be no conflict in the provision of network services between the local stub master and the original stack master. It should be noted that, according to various aspects of the current invention, instead of the backup configuration, no configuration may be applied to the local stub master in S180, and the local stub master may keep its original configuration.

[32] Fig. 4 presents an exemplary system diagram of various hardware components and other features, for use in accordance with aspects of the present invention. The present invention may be implemented using hardware, software, or a combination thereof and may be implemented in one or more computer systems or other processing systems. According to one aspect, the invention is directed toward one or more computer systems capable of carrying out the functionality described herein. An example of such a computer system 900 is shown in Fig. 4.

[33] Computer system 900 includes one or more processors, such as processor 904. The processor 904 is connected to a communication infrastructure 906 (e.g., a communications bus, cross-over bar, or network). Various software variations are described in terms of this exemplary computer system. After reading this description, it will become apparent to a person skilled in the relevant art(s) how to implement the invention using other computer systems and/or architectures.

[34] Computer system 900 can include a display interface 902 that forwards graphics, text, and other data from the communication infrastructure 906 (or from a frame buffer not shown) for display on a display unit 930. Computer system 900 also includes a main memory 908, preferably random access memory (RAM), and may also include a secondary memory 910. The secondary memory 910 may include, for example, a hard disk drive 912 and/or a removable storage drive 914, representing a floppy disk drive, a magnetic tape drive, an optical disk drive, *etc.* The removable storage drive 914 reads from and/or writes to a removable storage unit 918 in a well-known manner. Removable storage unit 918, represents a floppy disk, magnetic tape, optical disk, *etc.*, which is read by and written to removable storage drive 914. As will be appreciated, the removable storage unit 918 includes a computer usable storage medium having stored therein computer software and/or data.

[35] In accordance with aspects of the present invention, secondary memory 910 may include other similar devices for allowing computer programs or other instructions to be loaded into computer system 900. Such devices may include, for example, a removable storage unit 922 and an interface 920. Examples of such may include a program cartridge and cartridge interface (such as that found in video game devices), a removable memory chip (such as an erasable programmable read

only memory (EPROM), or programmable read only memory (PROM)) and associated socket, and other removable storage units 922 and interfaces 920, which allow software and data to be transferred from the removable storage unit 922 to computer system 900.

[36] Computer system 900 may also include a communications interface 924. Communications interface 924 allows software and data to be transferred between computer system 900 and external devices. Examples of communications interface 924 may include a modem, a network interface (such as an Ethernet card), a communications port, a Personal Computer Memory Card International Association (PCMCIA) slot and card, *etc.* Software and data transferred via communications interface 924 are in the form of signals 928, which may be electronic, electromagnetic, optical or other signals capable of being received by communications interface 924. These signals 928 are provided to communications interface 924 via a communications path (*e.g.*, channel) 926. This path 926 carries signals 928 and may be implemented using wire or cable, fiber optics, a telephone line, a cellular link, a radio frequency (RF) link and/or other communications channels. In this document, the terms “computer program medium” and “computer usable medium” are used to refer generally to media such as a removable storage drive 980, a hard disk installed in hard disk drive 970, and signals 928. These computer program products provide software to the computer system 900. The invention is directed to such computer program products.

[37] Computer programs (also referred to as computer control logic) are stored in main memory 908 and/or secondary memory 910. Computer programs may also be received via communications interface 924. Such computer programs, when executed, enable the computer system 900 to perform the features of the present invention, as discussed herein. In particular, the computer programs, when executed, enable the processor 910 to perform the features of the present invention. Accordingly, such computer programs represent controllers of the computer system 900.

[38] In a variation where the present invention is implemented using software, the software may be stored in a computer program product and loaded into computer

system 900 using removable storage drive 914, hard drive 912, or communications interface 920. The control logic (software), when executed by the processor 904, causes the processor 904 to perform the functions of the present invention as described herein. In another variation, the invention is implemented primarily in hardware using, for example, hardware components, such as application specific integrated circuits (ASICs). Implementation of the hardware state machine so as to perform the functions described herein will be apparent to persons skilled in the relevant art(s).

[39] In yet another variation, the present invention is implemented using a combination of both hardware and software.

[40] FIG. 5 is a block diagram of various exemplary system components, in accordance with aspects of the present invention. FIG. 5 shows a communication system 1000 usable in accordance with the present invention. The communication system 1000 includes one or more accessors 1060, 1062 (also referred to interchangeably herein as one or more "users") and one or more terminals 1042, 1066. In one embodiment, data for use in accordance with the present invention is, for example, input and/or accessed by accessors 1060, 1064 via terminals 1042, 1066, such as personal computers (PCs), minicomputers, mainframe computers, microcomputers, telephonic devices, or wireless devices, such as personal digital assistants ("PDAs") or a hand-held wireless devices coupled to a server 1043, such as a PC, minicomputer, mainframe computer, microcomputer, or other device having a processor and a repository for data and/or connection to a repository for data, via, for example, a network 1044, such as the Internet or an intranet, and couplings 1045, 1046, 1064. The couplings 1045, 1046, 1064 include, for example, wired, wireless, or fiberoptic links. One aspect of the current invention operates in a stand-alone environment, such as on a single terminal.

[41] While this invention has been described in conjunction with the exemplary aspects outlined above, various alternatives, modifications, variations, improvements, and/or substantial equivalents, whether known or that are or may be presently unforeseen, may become apparent to those having at least ordinary skill in the art. Accordingly, the exemplary aspects of the invention, as set forth above, are intended to

be illustrative, not limiting. Various changes may be made without departing from the spirit and scope of the invention. Therefore, the invention is intended to embrace all known or later-developed alternatives, modifications, variations, improvements, and/or substantial equivalents.

WHAT IS CLAIMED IS:

1. An apparatus for preventing network conflict in a stack, comprising:
a master node connected to a plurality of slave nodes via a network data link; and
a resiliency link connecting the master node to each one of the plurality of slave nodes; wherein
the resiliency link has at least one of control information only and a bandwidth that is lower than or equal to a bandwidth of the network link.
2. The apparatus of claim 1, wherein the stack is in at least one of a ring node configuration, a star configuration, a dual star configuration, and a mesh configuration.
3. The apparatus of claim 1, wherein at least one of the master node and the slave nodes store configuration files that permit a slave node to reconfigure into an alternative master node.
4. The apparatus of claim 1, wherein the resiliency link connects the master node to the plurality of slave nodes via a hub.
5. The apparatus of claim 4, wherein the hub is connected via the resiliency link to the master node and to each one of the slave nodes.
6. A method of preventing network conflict upon failure in a stack that includes a master node connected to a plurality of slave nodes via a network data link, the master node being configured by an original configuration file, the method comprising:
determining that a stack master is inoperative when no health-check notification has been received by at least one slave node within a period of time;

providing the health-check notification of a health of the master node to each slave node via a resiliency link; and

determining an appropriate configuration to apply upon a determination of stack master failure.

7. The method of claim 6, wherein the health-check notification is received via the resiliency link by the slave nodes.

8. The method of claim 6, further comprising:
reconfiguring one or more slave nodes using a fallback configuration file when the health-check notification indicates that the master node is operational.

9. The method of claim 6, further comprising:
reconfiguring one or more slave nodes using the original configuration file when a lapse of health-check notification indicates that the master node is not operational.

10. A system for preventing network conflict upon failure in a stack that includes a master node connected to a plurality of slave nodes via a network data link, the master node being configured by an original configuration file, the system comprising:

a determining module for determining that a stack master failure has occurred when no health-check notification has been received by at least one slave node within a predetermined period of time; and

a providing module for providing the health-check notification of a health of the master node to each slave node via a resiliency link; and

a determining module for determining an appropriate configuration to apply upon a determination of stack master failure.

11. The system of claim 10, wherein the providing module for providing a notification of the health of the master node to each slave node comprises a hub.

12. The system of claim 10, further comprising:
a reconfiguring module for reconfiguring one or more slave nodes using a fallback configuration file when the health-check notification indicates that the master node is operational.

13. The system of claim 12, wherein the reconfiguring module comprises a fallback configuration file.

14. A system for preventing network conflict upon failure in a stack that includes a master node connected to a plurality of slave nodes via a network data link, the master node being configured by an original configuration file, the system comprising:

a processor;
a user interface functioning via the processor; and
a repository accessible by the processor; wherein
a determination is made that a stack master failure has occurred when no health-check notification has been received by at least one slave node within a predetermined period of time; and
a notification of a health of the master node is provided to each slave node via a resiliency link; and
an appropriate configuration is applied upon a determination of stack master failure.

15. The system of claim 14, wherein the processor is housed on a terminal.

16. The system of claim 14, wherein the processor is housed on a server.

17. The system of claim 16, wherein the server is coupled to a network.

18. A computer program product comprising a computer usable medium having control logic stored therein for causing a computer to prevent network conflict upon failure in a stack that includes a master node connected to a plurality of slave nodes via a network data link and configured by an original configuration file, the control logic comprising:

first computer readable program code means for determining that a stack master failure has occurred when no health-check notification has been received by at least one slave node within a predetermined period of time; and

second computer readable program code means for providing a notification of a health of the master node to each slave node via a resiliency link; and

third computer readable program code means for determining an appropriate configuration to apply upon a determination of stack master failure.

19. The method of claim 6, wherein determining an appropriate configuration comprises precluding an application of a conflicting configuration on the plurality of slave nodes when the notification indicates that the master node is operational.

20. The system of claim 10, wherein the determining module for determining an appropriate configuration comprises a precluding module for precluding an application of a conflicting configuration on the plurality of slave nodes when the notification indicates that the master node is operational.

21. The system of claim 14, wherein the appropriate configuration being applied comprises a conflicting configuration being precluded from being applied among the plurality of slave nodes when the health notification indicates that the master node is operational.




Figure 1A




Figure 1B




Figure 1C




Figure 1D




Figure 2A




Figure 2B




Figure 2C

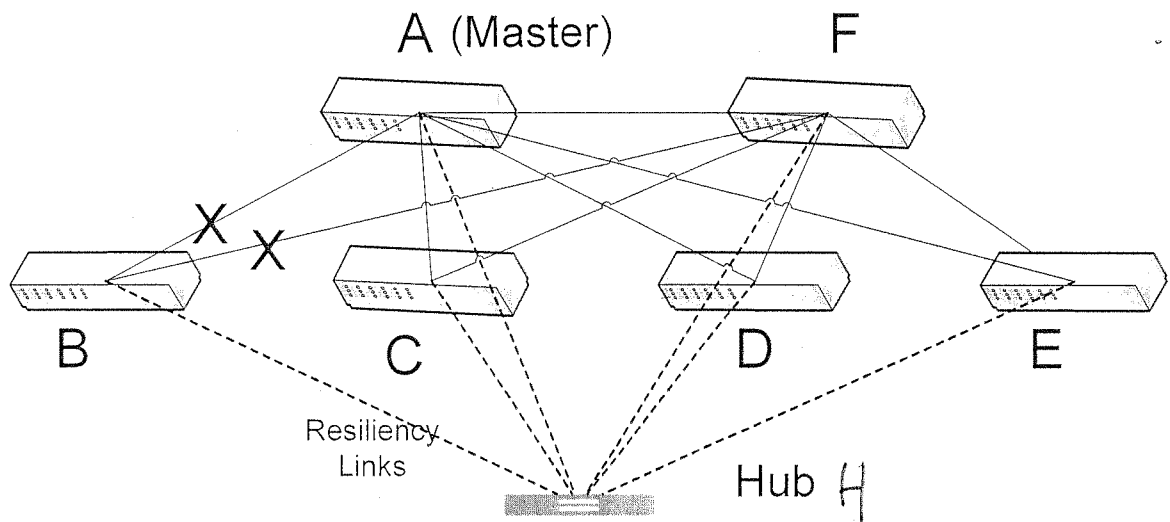


Figure 2D




Figure 2E

FIGURE 3

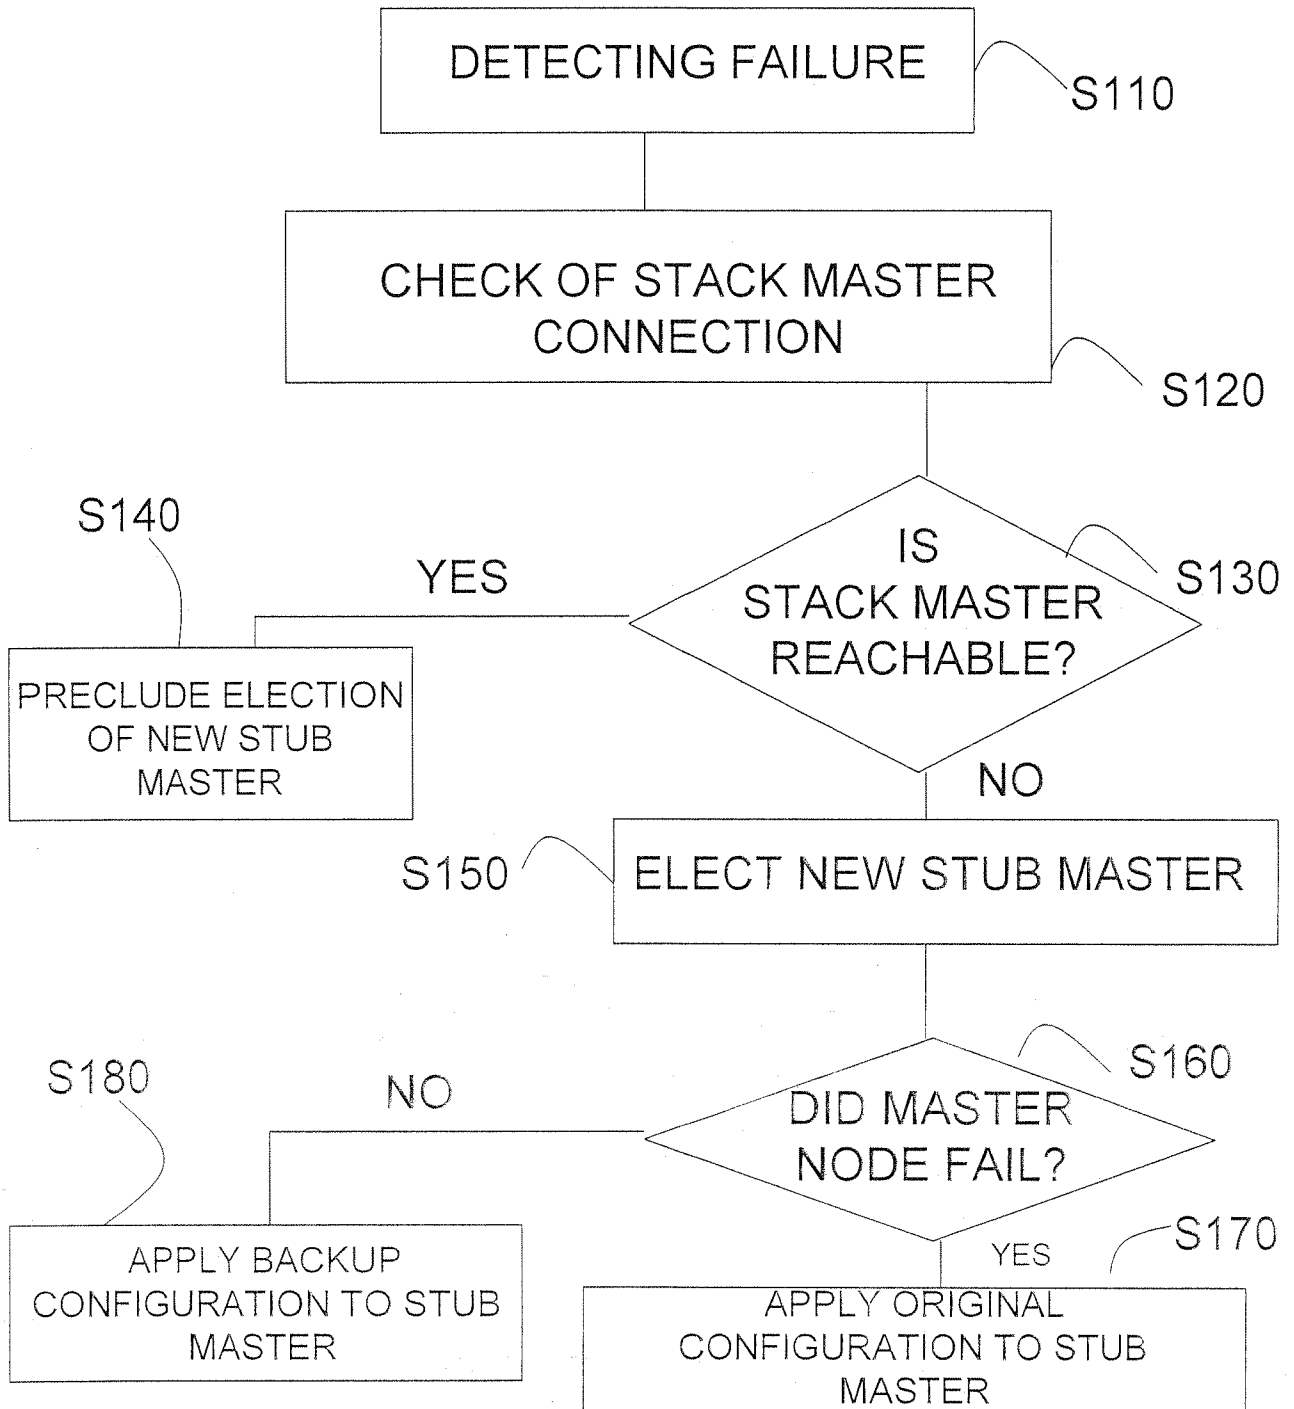


FIG. 4





FIG. 5



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 08/77838

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - H04L 12/403 (2008. 04)

USPC - 370/455

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)

IPC (8): H04L 12/403 (2008. 04)

USPC: 370/455

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 370/254, 431, 455, 461, 462; 709/201, 208, 209, 223, 226

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

Delphion; Google Scholar

Search Terms: network, redundancy, heartbeat, reconfigure, alternative, master, slave etc.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6,128,277 A (Bruck et al.) 03 October 2000 (03.10.2000), entire document, especially col 2 in 9-13, Fig. 1, col 2 in 20-25, col 10 in 64- col 11 in 6, Fig. 10A, col 6 in 20-22, Fig. 12, col 4 in 18-36, col 3 in 53-65, col 11 in 26-34, col 13 in 23-39, col 4 in 18-42, col 4 in 46-58, col 6 in 23-31.	1-21
A	US 2005/0144320 A1 (Jagana et al.) 30 June 2005 (30.06.2005)	1-21
A	US 2004/0081122 A1 (Koodli et al.) 29 April 2004 (29.04.2004)	1-21

Further documents are listed in the continuation of Box C.

* 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

16 November 2008 (16.11.2008).

Date of mailing of the international search report

02 DEC 2008

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-6300
PCT OSP: 571-272-7774