

5 FAULT-TOLERANT COMMUNICATIONS IN ROUTED NETWORKS

BACKGROUND

[0001] In a computer networking environment, multiple nodes may communicate with each other over a network. Should the network experience a failure, communication between the nodes may be disrupted.

10 SUMMARY

[0002] The following presents a simplified summary of the disclosure in order to provide a basic understanding to the reader. This summary is not an extensive overview of the disclosure and it does not identify key or critical elements of the invention or delineate the scope of the invention. Its sole purpose is to
15 present some concepts disclosed herein in a simplified form as a prelude to the more detailed description that is presented later.

[0003] The following examples provide computer network communication fault-tolerance via unique network stack architectures requiring minimal consideration by application software operating on networked nodes.

20 [0004] Many of the attendant features will be more readily appreciated as they become better understood by reference to the following detailed description considered in connection with the accompanying drawings.

DESCRIPTION OF THE DRAWINGS

[0005] The present description will be better understood from the following
25 detailed description read in light of the accompanying drawings, wherein:

[0006] FIG. 1 is block diagram showing an example network stack architecture.

[0007] FIG. 2 is a block diagram showing a networked computing environment including two example nodes coupled via two networks.

30 [0008] FIG. 3 is a block diagram showing an example fault-tolerant communications driver, NETFT.

[0009] FIG. 4 is a block diagram showing an example fault-tolerant communications architecture including NETFT and an application.

5 [0010] FIG. 5 is a flow diagram showing data flowing through a fault-tolerant communications environment including a source node and a destination node coupled via Path A over network 1 and Path B over network 2.

[0011] FIG. 6 is a flow diagram showing data flowing through the fault-tolerant communications environment shown in FIG. 5 with the addition of several
10 possible communications failures.

[0012] FIG. 7 is a block diagram showing and another example of a fault-tolerant communications driver, NETFT.

[0013] FIG. 8 is a block diagram showing an example fault-tolerant communications architecture including NETFT and an application.

15 [0014] FIG. 9 is a flow diagram showing data flowing through a fault-tolerant communications environment including a source node and a destination node coupled via Path A over network 1 and Path B over network 2.

[0015] FIG. 10 is a flow diagram showing data flowing through the fault-tolerant communications environment shown in FIG. 9 with the addition of several
20 possible communications failures.

[0016] FIG. 11 is a block diagram showing an example computing environment in which the technology described above may be implemented.

[0017] Like reference numerals are used to designate like parts in the accompanying drawings.

25 DETAILED DESCRIPTION

[0018] The detailed description provided below in connection with the appended drawings is intended as a description of the present examples and is not intended to represent the only forms in which the present examples may be constructed or utilized. The description sets forth the functions of the examples and
30 the sequence of steps for constructing and operating the examples. However, the same or equivalent functions and sequences may be accomplished by different examples.

[0019] Although the present examples are described and illustrated herein as being implemented in a computing and networking system, the system described is

5 provided as an example and not a limitation. As those skilled in the art will appreciate, the present examples are suitable for application in a variety of different types of computing and networking systems.

[0020] FIG. 1 is block diagram showing an example network stack architecture 100. A network stack ("stack") generally couples with software applications via network stack interfaces and/or other interfaces to provide network communications functionality to applications. An application is typically said to be at (or coupled to) the "top" of the stack. A network is typically said to be at (or coupled to) the "bottom" of the stack. Various elements of a network stack may be referred to as at or near the top or bottom of the stack, or higher or lower in the stack relative to each other. For example, in FIG.1, protocol driver 130 is higher in the stack than NIC 180 which is shown at the bottom of the stack in this particular figure. Various depictions of a network stack may or may not include some stack elements, or may group, order or name the elements in various ways, depending on the purpose or focus of the depiction, as understood by those skilled in the art.

20 [0021] The term "driver" as used herein refers to a control program or the like that enables a node to operate with a particular device, such as a printer, network interface card, or other computer subsystem, or to operate with one or more programs such as network stacks, protocol drivers, and/or other computer software or firmware or the like. For example, a protocol driver typically operates with a network stack.

[0022] An application may pass a packet of data to a stack destined for an application operating on another node. In this case, the data is said to flow "down" the stack and is sent out over a network. Data received by a node is said to flow "up" the stack until it reaches the destined application. Such networking systems are well known to those skilled in the art.

30 [0023] In one example a stack is based on the Network Driver Interface Specification ("NDIS") which defines a standard application programming interface ("API") for network interface cards ("NICs"), such as NIC 180, and abstracts the network hardware from network drivers. NDIS also specifies a standard interface

5 between layered network drivers, thereby abstracting lower-level drivers that manage hardware, such as a miniport driver, from upper-level drivers, such as protocol drivers. Multiple NDIS-conforming protocol drivers may co-exist on a single node. Also, if a node includes multiple NICs, perhaps because it is connected to more than one network, NDIS routes network traffic to the appropriate NIC via its
10 associated driver as indicated by the traffic. An illustration of NDIS is shown in FIG. 1. Other networking stack standards, technologies and/or architectures, such as the Open Data-Link Interface ("ODI"), the Data Link Provider Interface ("DLPI"), the Uniform Driver Interface ("UDI"), or other technologies, may be used with the following examples as well with appropriate modifications, as would be understood
15 by those skilled in the art. As a matter of convenience, NDIS and NDIS terminology is used with examples throughout this description, but other standards, technologies and/or architectures may be used in all of these examples with appropriate modifications, unless otherwise noted.

[0024] As shown in FIG. 1, coupled to NIC 180 via NDIS 120 is miniport driver
20 160. A miniport driver typically interacts with NDIS via an NDIS miniport interface 162. The miniport driver 160 may be associated with NIC 180 and may manage its operations, including sending and receiving data through the NIC. The miniport driver 160 typically interfaces with higher-level drivers, such as intermediate driver 140 and protocol driver 130. A miniport driver is considered a NIC driver. NIC
25 miniports generally perform those hardware-specific operations needed to manage a particular NIC with common or NIC-independent functionality provided by NDIS. A node may include multiple NICs with each NIC generally having an associated NIC driver. Some examples in this description describe the use of miniport drivers but, as will be understood by those skilled in the art, any type of NIC driver or the like may
30 be used in these examples, unless otherwise noted.

[0025] Protocol or transport driver 130 couples to NDIS 120 via an NDIS protocol interface 134. Protocol drivers or transport protocol drivers generally provide the functionality to create, send and receive packets of data that are sent from one node to another through the network stack and over a network. As known

5 to those skilled in the art, a common reliable or guaranteed delivery transport protocol may be TCP/IP (Transmission Control Protocol / Internet Protocol). UDP (User Datagram Protocol) over IP may be a common unreliable or non-guaranteed delivery protocol. TCP, UDP and/or other protocols, such as IPX/SPX (Internet Packet Exchange / Sequenced Packet Exchange), may be used with the following examples
10 unless otherwise noted.

[0026] NDIS intermediate ("IM") drivers 140 are shown between protocol drivers 130 and NDIS NIC miniports 160 in FIG. 1. To protocol drivers IM drivers appear to be NDIS miniports while to NIC drivers they look like protocol drivers. Data packets flowing up or down the network stack pass through the IM driver 140 which
15 may ignore, inspect, filter, forward, redirect and/or modify the data packets. An intermediate driver 140 may also be known as a filter driver.

[0027] FIG. 2 is a block diagram showing a networked computing environment 200 including two example nodes 210 and 260 coupled via two networks 202 and 282. Nodes 210 and 260 may each be personal computers ("PCs"),
20 client computers, servers, hosts, laptops, portable devices, consumer electronic devices, or any of various other types of computing or processing devices, machines or systems. One non-limiting example of a type of computing system is described in detail below with respect to FIG. 11. Circles 212, 214, 262, and 264 represent NICs associated with their respective nodes. One non-limiting example of a type of NIC is
25 further described below with respect to FIG. 11 as network adapter 1113.

[0028] As used herein, the term node refers to any computing system, device, or process that is uniquely addressable, or otherwise uniquely identifiable, in a network (e.g., network 202) and that is operable to communicate with other nodes in the network. For example, and without limitation, a node may be a personal
30 computer, a server computer, a hand-held or laptop device, a tablet device, a multiprocessor system, a microprocessor-based system, a set top box, a consumer electronic device, a network PC, a minicomputer, a mainframe computer, or the like. A non-limiting example of a node 210, in the form of a computing system, is set forth below with respect to FIG. 11.

- 5 [0029] Networks 202 and 282 may be the same network, may exist on the same or different subnets, may be logically or physically coupled or isolated from each other, may use similar or different networking technologies, etc. In particular, networks 202 and 282 may be routed networks, that is, networks including routers that forward routable protocol packets. Routable protocols are typically considered
- 10 communications protocols used to route data from one network to another. An example of a routable protocol is TCP/IP. Sending a data packet in a routable fashion implies using a routable transport protocol to format and/or send the data packet. Those skilled in the art will be familiar with routable protocols and routing network topologies, systems and architectures.
- 15 [0030] In one example, networks 202 and 282 may be independent of each other such that if there is a problem or failure with one network it is unlikely to affect the operational status of the other. In other examples, three or more networks may be used. In examples where greater degrees of fault-tolerance are desired a larger number of networks along with the associated connectivity of nodes to those
- 20 networks, including a similar number of NICs installed on a node, may be employed.
- [0031] NIC 212, associated with node 210, is shown with an example address of 172.56.48.37 and is coupled to network 1 202. NIC 214, also associated with node 210, is shown with an example address of 197.71.48.38 and is coupled to network 2 282. NIC 262, associated with node 260, is shown with an example
- 25 address of 172.56.48.38 and is also coupled to network 1 202. NIC 264, also associated with node 260, is shown with an example address of 197.71.48.39 and is also coupled to network 2 282. These addresses may, in practice, be IPv4 or IPv6 addresses or the like, or any other type of network address typically related to the protocol being used.
- 30 [0032] Each node may include one or more NICs. Arrows 201 and 203, also shown in FIG. 11 as arrow 1114, represent a first communications route or pathway ("Path A") over network 1 202 between nodes 210 and 260. Arrows 281 and 283 represent a second communications route or pathway ("Path B") over network 2 282 between nodes 210 and 260. In practice, there may be one or more pathways over

5 one or more networks between the two or more nodes in environment 200. The term “pathway” as used herein is defined as a communications route, or communications link, between nodes in a network. Such a route or link may be dynamic in that the exact route between nodes may change over time.

[0033] Blocks 216 and 266 represent an application and a network stack,
10 including a fault-tolerant communications (“FT”) driver, provided on each of nodes 210 and 260. The FT driver of block 216 is shown with an example address of 10.0.0.1 and the FT driver of block 266 is shown with an example address of 10.0.0.2. These addresses are typically considered virtual addresses. These addresses may be IPv4 or IPv6 addresses or the like, or any other type of network or
15 communications address. FT drivers may or may not have virtual addresses as shown in the various examples below.

[0034] A fault-tolerant network stack is a network stack including an FT driver, such as NETFT described below in connection with FIG. 3, or the like. An FT driver, such as NETFT, operating in combination with a network stack typically allows
20 nodes to communicate with each other via one or more communications paths, such as Path A and Path B, over one or more networks. Should any of these communications paths fail, the nodes may continue communicating given at least one operational pathway. Such a pathway failure may result from failure of a NIC or failure of any element of a pathway, including connections, cabling or other
25 communications media (including radio frequency (“RF”) or infrared (“IR”) and the like), routers, hubs, switches, firewalls, Internet Service Providers (“ISPs”), power failure to any node, device or system of the network, or the like.

[0035] In one example, a communications failure may result in a plug-and-play (“PnP”) event. A PnP event may indicate the removal of a NIC from its node or to
30 a media sense change. A media sense disconnect, for example, typically results from a failure that causes the NIC to lose the signal or carrier on the network media, such as a network cable, RF or IR link or the like. A media sense disconnect may be caused by disconnecting the network cable or carrier from the NIC or powering off the other end of the cable (a hub or switch, for example). A media sense connect is typically

5 the opposite, such as reconnecting the cable, re-powering on the hub or switch or the like. These types of events, also known as connectivity events, are generally local events in that they occur on or are proximate to the node itself. Such local connectivity events typically result in an event indication, such as a PnP event or the like, on a node.

10 [0036] In another example, a communications failure may be detected by using heartbeat packets sent between nodes. Failure of such a heartbeat packet may indicate failure of a pathway between nodes. Heartbeat packets tend to be marked such that the FT driver can detect them upon receipt and remove them for the packet flow being passed up the network stack. In one example, heartbeat packets may be
15 implemented using route control protocol ("RCP") by forming RCP packets. Such heartbeat packets may be used to validate the end-to-end operational status of a pathway. That is, by sending a heartbeat packet from node 210 over Path A to node 260 and by node 210 receiving a reply to the sent heartbeat packet from node 260, it is generally considered that Path A is end-to-end operational. Should the
20 heartbeat fail (no heartbeat reply received in response to the heartbeat sent), such a failure may indicate that Path A is not operational, perhaps due to failure of some element of network 1 202 such as a router, switch, connection, or the like, or due to the target node itself failing. In particular, node 210 may have an operational NIC 212 and valid media sense, indicating that it is properly connected to the network,
25 but may still detect a heartbeat failure due to some network or system failure down the line.

[0037] FIG. 3 is a block diagram showing an example fault-tolerant communications driver, NETFT 300. NETFT 300 may be implemented as an NDIS miniport driver (FIG. 1, 160) for use with an NDIS network stack and for providing
30 network communications between nodes tolerant of pathway failures. That is, communications between two or more nodes may continue when each is using a NETFT despite failure of any component in the pathway as long as at least one pathway remains operational.

5 [0038] In one example, implementation of the FT driver as an NDIS miniport driver provides at least two benefits. First, because such an FT driver generally sits below any protocol drivers in the stack, protocol reliability tends to be provided by any higher-level reliable protocol driver which is generally unaffected by the addition of link-level fault-tolerance provided by an FT driver. For example, when using an FT
10 driver in combination with a protocol driver such as a TCP/IP driver, the FT driver will typically detect failed pathways and route data packets over end-to-end operational pathways independent of any protocol driver. Should any packet loss occur due to switching pathways, the TCP/IP protocol driver, which generally sits above the FT driver in the stack, tends to detect such losses and perform any retry or resend
15 operations to ensure that the reliable protocol succeeds in packet delivery.

[0039] A second benefit of placing the FT driver below the protocol driver in the stack is that typically no degradation of the routability of the protocol is introduced. When so configured, any tunneling operation that an FT driver performs on a data packet may employ a routable protocol, such as TCP or UDP, thus ensuring
20 that such data is routable, in addition to being link-level fault tolerant. To "routeably tunnel" a data packet is to tunnel a data packet using a routable protocol.

[0040] NETFT, as a part of a network stack, generally couples to a software application via NDIS or other network stack interfaces. Such a coupling generally enables applications to send and receive data packets over networks coupled to the
25 bottom of the stack. In one example, applications tend to use a virtual address as the source address for their data packets, this virtual address being known to NETFT and mapped and communicated to other nodes on the network as described below. As shown in FIG. 3, NETFT includes a miniport adapter 302 (also known as a processing element) a routing database 304, and one or more route monitor
30 adapters 306 and tunnel adapters 308.

[0041] Tunnel adapter 308 typically represents one NIC on the local node (or, in some instances, a virtual NIC) and maintains a socket used to tunnel packets to NETFT on the target node. There is typically one tunnel adapter 308 associated with each NIC on the local node with each NIC being coupled to a network providing a

5 pathway to another node. Each network may or may not be isolated from any other network. A tunnel adapter 308 is typically associated with a tunneling protocol driver and tunnels data packets through a tunneling protocol to and from its associated NIC via NDIS interfaces. One example of a tunneling protocol is UDP. Alternatively, other protocols, such as TCP, IPX, or SPX, may be used for tunneling. A tunnel
10 adapter 308 may become inactive should the associated NIC or media connection become inactive.

[0042] A routing database 304, as implemented in NETFT, is typically a simple data structure, that may be located in system memory, that includes entries mapping a virtual address for one or more pathways to a similar NETFT on another
15 node. In one example, mappings are represented by route monitor adapters such as route monitor adapter 306 which are typically associated with a tunnel adapter such as tunnel adapter 308. Generally a routing database such as routing database 304 will include one set of route adapters for each tunnel adapter, each route adapter being associated with a different target node reachable over the pathway associated
20 with the tunnel adapter. When using TCP/IP, for example, the database may map a destination virtual address to a physical address of a specific remote node.

[0043] A routing database 304 may also include priority information for each pathway. Such priority information may be used to indicate a preferred or primary pathway to another node and/or may include information about pathway speed or
25 other characteristics. A preferred pathway is the pathway calculated by NETFT to be used over other possible pathways, when possible, based on priority information and/or pathway status. Priority information may alternatively indicate a round-robin load balancing algorithm for making use of multiple pathways to a target node to load-balance traffic over the pathways, or enable some other pathway prioritization
30 scheme.

- 5 [0044] An example routing table database 304 mapping table is shown in Table 1.

Destination Address Type	Address	Priority
Virtual	10.0.0.2	--
Physical: Path A	172.56.48.38	1
Physical: Path B	197.71.48.39	2

Table 1

- 10 [0045] Referring to table 1 and FIG. 2, table 1 shows an example mapping table as might be used by NETFT operating on node 216. Table 1 shows virtual destination address 10.0.0.2, the virtual address as shown for node 266, mapped to physical address 172.56.48.38 associated with Path A to node 266 and physical address 197.71.48.39 associated with Path B to node 266. Path A is shown with first
 15 priority and Path B with second priority. Table 1 is provided as an example and is not intended to be limiting.

- [0046] When sending data from node 216 to node 266, such a mapping table is typically used to tunnel a packet destined to virtual destination address 10.0.0.2 by forwarding the packet via a tunneling protocol, such as UDP, to physical
 20 destination address 172.56.48.38, thus tunneling the packet from node 216 over Path A to node 266. One such mapping table may be created in the routing database (FIG.3, 304) for each set of pathways established between two nodes. Such a mapping table may be implemented in various forms, use various priority schemes and/or store other information including pathway operational status. The mapping
 25 table structure, number of pathways, address formats, etc. shown in Table 1 are provided as examples and are not intended to be limiting.

- [0047] The local node virtual address, remote node virtual addresses, and priority and other pathway information are typically provided to nodes by an out-of-band mechanism and passed to NETFT via its NDIS interfaces. This out-of-band
 30 mechanism may be as simple as a systems administrator using a management

5 application to specify the information, or it may be an automated system or the like. Such out-of-band mechanisms are well known to those skilled in the art.

[0048] As shown in FIG. 3, miniport adapter 302 (also known as the processing element of the driver) typically parses a data packet flowing down the network stack, examines the destination virtual address of the packet and uses
10 information from the routing database 304 to determine which tunnel adapter 308 to tunnel the data packet through. Incoming packets, or data packets flowing up the stack, are forwarded up the stack toward their destination virtual address, the tunneling protocol having previously removed the tunneling packet headers. In particular, the tunnel adapter 308 inspects incoming packets and forwards heartbeat
15 packets to a route monitor adapter 306 and forwards other packets up the stack via a miniport adapter 302. Aspects of tunneling data packets using a tunneling protocol and how protocol headers are added and removed by protocol drivers is well known to those skilled in the art.

[0049] Route monitor adapter 306 typically represents a remote node
20 accessible over a specific pathway identified by an associated tunnel adapter. The route monitor adapter 306 will typically provide a physical address for the remote node, the physical address also corresponding to a specific pathway to the remote node. This physical address is typically used for mappings in a routing database 304. There is typically one route monitor adapter for each distinct pathway to a remote
25 node, each route monitor adapter being associated with a tunnel adapter representing a pathway. In one example, referring back to FIG. 2, node 210 is shown coupled to node 260 over two pathways, one through network 1 202 ("Path A") and the other through network 2 282 ("Path B"). NETFT operating on node 210 may include a first route monitor adapter ("RMA-A") providing remote node 260's
30 physical address 172.56.48.38 associated with its NIC 262. RMA-A may be associated with a first tunnel adapter ("TA-A") on node 210 which may be associated with Path A. NETFT on node 210 may also include a second route monitor adapter ("RMA-B") providing remote node 260's second physical address 197.71.48.39

5 associated with its NIC 264. RMA-B may be associated with a second tunnel adapter ("TA-B") on node 210 which may be associated with Path B.

[0050] Referring to FIG. 3, route monitor adapter 306 typically monitors the health of a pathway to a remote node and indicates a failed or non-operational pathway in the routing database 304. Monitoring typically includes receiving any
10 event indications and/or noting any heartbeat failures and updating the database 304 accordingly. In one example, an event indicating the failure of a NIC or media connection may result in the disabling of the tunnel adapter 308. In another example, a heartbeat failure may result in the disabling of the route monitor adapter 306 associated with the specific remote node for which the heartbeat failed.

15 [0051] FIG. 4 is a block diagram showing an example fault-tolerant communications architecture 216 including NETFT 300 and an application 402. In this example, the application 402 sends data packets to NETFT 300 via the stack using a virtual source address 217 and a virtual destination address representing the destination node. Such out-going data packets flow via path 480 from the
20 application and through the network stack to the driver 300. The driver 300 typically determines which of the possible pathways each packet should take, generally using priority information and pathway operational status information stored in the routing database, and tunnels the packet to the target node over the selected pathway using the appropriate physical source address 422 or 424.

25 [0052] Application 402 may send a data packet through NETFT 300 via the TCP protocol, as shown in FIG. 4. Alternatively UDP or any other protocol may be used. Also, as shown, NETFT 300 may use the UDP protocol to tunnel packets to the target node. Alternatively, TCP or any other protocol may be used for tunneling. Further, alternate examples may not make use of miniport adapters or NDIS drivers
30 but may use other mechanisms or architectures to perform similar functions. Finally, the various elements of the network stack and the like may operate in either a User Mode or a Kernel Mode, either as shown or otherwise, or on systems with or without equivalent modes of operation.

5 [0053] FIG. 5 is a flow diagram showing data flowing through a fault-tolerant communications environment 500 including a source node 216 and a destination node 266 coupled via Path A over network 1 202 and Path B over network 2 282. In this example environment 500, data is shown being sent from the application operating on node 216 to the application listening on the destination virtual address
10 on node 266. The data packets flow down the network stack operating on node 216 using the TCP protocol into NETFT as shown by path 501. Assuming, as shown, that Path A is the selected pathway, NETFT maps the data packets from the source virtual address being used by the application to Path A and tunnels the data through the UDP protocol using the Path A physical destination address for target node 266, out
15 NIC 1 of node 216 as further shown by path 501 and onto network 1 202 via link 201. The data then flows through network 1 202, over link 203 and to node 266, flowing up the network stack operating on node 266 as shown by path 503. The data then flows through the UDP protocol driver, the same protocol that was used on the sending side as the tunneling protocol, where the UDP protocol headers are stripped
20 off the data packets which are then passed into NETFT operating on node 266. NETFT then forwards the data packets up the stack to the application which is listening on the destination virtual address. Responses tend to flow in the reverse order.

[0054] FIG. 6 is a flow diagram showing data flowing through the fault-tolerant communications environment 500 shown in FIG. 5 with the addition of
25 several possible communications failures 610, 612, 620, 622, and 630. Other communications failures are also possible. Failure 610 indicates a failure of NIC 1 operating on the sending node 216. Such a failure may occur should NIC 1 be removed from the node, should the driver of NIC 1 fail, should NIC 1 itself fail, or the
30 like. The failure may be detected by NETFT via an event indication, such as a PnP event or the like, and/or a heartbeat failure. In such a situation Path A is typically considered to have failed and NETFT will select an alternate end-to-end-operational pathway. An end-to-end operational pathway is typically a pathway that can

5 successfully deliver data from the source node and application all the way to the destination node and application.

[0055] Failure 620 indicates a failure of the network media coupling with NIC1 of node 216. This failure may be due to a cable being disconnected from NIC 1, from the cable becoming disconnected from some device of network 1 202, from the
10 device the cable is connected to on the network side being powered down or failing, or the like. This type of failure may also be detected by NETFT via an event indication, such as a PnP event or the like, and/or a heartbeat failure and an alternate pathway selected.

[0056] Failure 630 indicates a failure of some type within network 202
15 resulting in data packets failing to reach destination node 266. In this failure case, sending node 216 may still be coupled to network 202 with a proper media sense indication, yet Path A has become disrupted further down the network. Given such a failure, NETFT operating on sending node 216 may not detect the failure via an event indication if local indications show connectivity to the network 202 as good, but may
20 detect the failure via Path A heartbeat failure.

[0057] Failure 622 of link 203 and failure 612 of NIC 1 operating on receiving node 266 tend to be similar to the corresponding failures shown for node 216. But these failures, not being local to node 216 may not be detected via event indications but may be detected via heartbeat failure.

25 [0058] Any of these failures, and other failures, may be detected by NETFT operating on node 216 and result in it selecting an alternate end-to-end operational pathway, such as Path B over network 2 282. In this example, as shown in FIG. 6, NETFT tunnels data down alternate path 681 and over network 2 282 to receiving node 266. Should the failure condition be corrected and end-to-end operational
30 status restored on Path A, NETFT operating on sending node 216 may detect the recovery and again make use of Path A. Further, any responses from node 266 back to node 216 may be tunneled in a similar fault-tolerant fashion by NETFT.

[0059] FIG. 7 is a block diagram showing and another example of a fault-tolerant communications driver, NETFT 700. This example is similar to the example

5 shown in FIG. 3 but includes variations as described below. In this example, a software application may not need to use virtual addresses. Instead, an application may use a physical destination address to address data packets to the target node.

[0060] Protocol adapter 710 generally couples to miniport adapter 702 (also known as the processing element of the driver) and to a NIC miniport adapter (not
10 shown). There is typically one protocol adapter for each NIC installed on the node, each protocol adapter being associated with a NIC via its NIC adapter. As each protocol adapter is associated with a NIC, it is also associated with the pathway coupled to the NIC. The protocol adapter 710 is operable to accept data packets from an application via the processing element 702 and pass the data packets to the
15 associated NIC without the need for tunneling.

[0061] Processing element 702 typically parses a data packet flowing down the network stack, examines the physical destination address of the packet and uses information from the routing database 704 to determine if the packet can be forwarded over a protocol adapter 710 or needs to be tunneled over a tunnel adapter
20 308 to the target node. Generally, if the pathway indicated by the physical destination address is end-to-end operational, the data packet will be sent over that pathway. Otherwise and alternate pathway may be selected over which the packet may be tunneled.

5 [0062] In this example the routing database 704 maintains mappings of physical destination addresses and pathways, along with priority and other information as described above. An example routing database 704 mapping table is shown in Table 2.

Destination Address Type	Address	Priority
Physical: Path A	172.56.48.38	1
Physical: Path B	197.71.48.39	2

10 Table 2

[0063] Referring to Table 2 and FIG. 2, Table 2 shows an example mapping table as might be used by NETFT operating on node 216. Table 2 shows a mapping including physical destination address 172.56.48.38 associated with Path A to node 15 266 and the physical destination address 197.71.48.39 associated with Path B to node 266. Path A is shown with first priority and Path B with second priority.

[0064] When sending data from node 216 to node 266, such a mapping table is typically used in forwarding (or tunneling if needed) a data packet being sent to physical destination address 172.56.48.38 of node 266. If the pathway associated 20 with the original destination address is operational, the data packet tends to be forwarded to the destination node without tunneling. If that pathway is not available, then the data packet is sent over the alternate pathway to physical destination address 197.71.48.39 of node 266 via tunneling. Other aspects of NETFT 700 are generally similar to those of NETFT as described for FIG. 3.

25 [0065] FIG. 8 is a block diagram showing an example fault-tolerant communications architecture 216 including NETFT 700 and an application 402. In this example, application 402 sends data packets to NETFT 700 via the stack using a physical source address and a physical destination address 801 representing the destination node. Such out-going data packets flow via path 880 from the 30 application and through the network stack to the driver 700. The driver 700 typically determines which of the possible pathways each packet should take, generally using

5 priority information and pathway operational status information stored in the routing database, and either forwards the packet to the target node over pathway indicated by the original physical destination address or, if that pathway is not end-to-end operation, tunnels the packet over an alternate pathway as indicated in this example by route 882 and NIC 2 892.

10 [0066] Application 402 may send a data packet through NETFT 700 via the TCP protocol, as shown in FIG. 8. Alternatively UDP or any other protocol may be used. Also, as shown, NETFT 700 may use the UDP protocol to tunnel packets to the target node. Alternatively, TCP or any other protocol may be used for tunneling. Further, other examples may not make use of NDIS drivers but may use other
15 mechanisms or architectures to perform similar functions. Finally, the various elements of the network stack and the like may operate in either a User Mode or a Kernel Mode, either as shown or otherwise, or on systems with or without equivalent modes of operation.

[0067] FIG. 9 is a flow diagram showing data flowing through a fault-tolerant
20 communications environment 900 including a source node 816 and a destination node 966 coupled via Path A over network 1 202 and Path B over network 2 282. In this example environment 900, data is shown being sent from the application operating on node 216 to the application listening on the destination physical address on node 266. The data packets flow down the network stack operating on
25 node 216 using the TCP protocol into NETFT as shown by path 901. Assuming, as shown, that Path A is the selected pathway, NETFT forwards the data packets using the physical destination address provided by the application via NIC 1 of node 216 over Path A and network 1 202 via link 201. The data then flows through network 1 202, over link 203 and to node 966, flowing up the network stack operating on node
30 966 as shown by path 903. The data then flows through NETFT and the protocol driver (a protocol driver for the same protocol that was used on the sending side as the sending protocol) and up to the application. Responses tend to flow in the reverse order.

5 [0068] FIG. 10 is a flow diagram showing data flowing through the fault-tolerant communications environment 900 shown in FIG. 9 with the addition of several possible communications failures 1010, 1012, 1020, 1022, and 1030. Other communications failures are also possible. Failure 1010 indicates a failure of NIC 1 operating on the sending node 816. Such a failure may occur should NIC 1 be removed from the node, its NIC driver fail, the NIC itself fail, or the like. The failure may be detected by NETFT via an event indication, such as a PnP event or the like, and/or a heartbeat failure. In such a situation Path A is typically considered to have failed and NETFT will select an alternate end-to-end-operational pathway.

[0069] Failure 1020 indicates a failure of the network media coupling with NIC1 of node 816. This failure may be due to a cable being disconnected from NIC 1, from the cable becoming disconnected from some device of network 1 202, from the device the cable is connected to on the network side being powered down or failing, or the like. This type of failure may also be detected by NETFT via an event indication, such as a PnP event or the like, and/or a heartbeat failure and an alternate pathway selected.

[0070] Failure 1030 indicates a failure of some type within network 202 resulting in data packets failing to reach destination node 966. In this failure case, sending node 816 may still be coupled to network 202 with a proper media sense indication, yet Path A has become disrupted further down the network. Given such a failure, NETFT operating on sending node 816 may not detect the failure via an event indication, such as a PnP event or the like, if local indications show connectivity to the network 202 as good, but may detect the failure via Path A heartbeat failure.

[0071] Failure 1022 of link 203 and failure 1012 of NIC 1 operating on receiving node 966 tend to be similar to the corresponding failures shown for node 816. But these failures, not being local to node 816 may not be detected via event indications but may be detected via heartbeat failure.

[0072] Any of these failures, and other failures, may be detected by NETFT operating on node 816 and result in it selecting an alternate end-to-end operational pathway, such as Path B over network 2 282. In this example, as shown in FIG. 10,

5 NETFT tunnels data down alternate path 1081 and over network 2 282 to receiving node 966. Should the failure condition be corrected and end-to-end operational status restored on Path A, NETFT operating on sending node 816 may detect the recovery and again make use of Path A. Further, any responses from node 966 back to node 816 may be forwarded or tunneled, depending on the operational status of
10 Path A and Path B, in a similar fault-tolerant fashion by its NETFT.

[0073] FIG. 11 is a block diagram showing an example computing environment 1100 in which the technology described above may be implemented. A suitable computing environment may be implemented with numerous general purpose or special purpose systems. Examples of well known systems may include,
15 but are not limited to, personal computers ("PC"), hand-held or laptop devices, microprocessor-based systems, multiprocessor systems, servers, workstations, consumer electronic devices, set-top boxes, and the like.

[0074] Computing environment 1100 generally includes a general-purpose computing system in the form of a computing device 1101 coupled to various
20 peripheral devices 1102, 1103, 1104 and the like. System 1100 may couple to various input devices 1103, including keyboards and pointing devices, such as a mouse or trackball, via one or more I/O interfaces 1112. The components of computing device 1101 may include one or more processors (including central processing units ("CPU"), graphics processing units ("GPU"), microprocessors ("uP"),
25 and the like) 1107, system memory 1109, and a system bus 1108 that typically couples the various components. Processor 1107 typically processes or executes various computer-executable instructions to control the operation of computing device 1101 and to communicate with other electronic and/or computing devices, systems or environment (not shown) via various communications connections such
30 as a network connection 1114 or the like. System bus 1108 represents any number of several types of bus structures, including a memory bus or memory controller, a peripheral bus, a serial bus, an accelerated graphics port, a processor or local bus using any of a variety of bus architectures, and the like.

5 [0075] System memory 1109 may include computer readable media in the form of volatile memory, such as random access memory ("RAM"), and/or non-volatile memory, such as read only memory ("ROM") or flash memory ("FLASH"). A basic input/output system ("BIOS") may be stored in non-volatile or the like. System memory 1109 typically stores data, computer-executable instructions and/or
10 program modules comprising computer-executable instructions that are immediately accessible to and/or presently operated on by one or more of the processors 1107.

[0076] Mass storage devices 1104 and 1110 may be coupled to computing device 1101 or incorporated into computing device 1101 via coupling to the system
15 bus. Such mass storage devices 1104 and 1110 may include a magnetic disk drive which reads from and/or writes to a removable, non-volatile magnetic disk (e.g., a "floppy disk") 1105, and/or an optical disk drive that reads from and/or writes to a non-volatile optical disk such as a CD ROM, DVD ROM 1106. Alternatively, a mass storage device, such as hard disk 1110, may include non-removable storage
20 medium. Other mass storage devices may include memory cards, memory sticks, tape storage devices, and the like.

[0077] Any number of computer programs, files, data structures, and the like may be stored on the hard disk 1110, other storage devices 1104, 1105, 1106 and system memory 1109 (typically limited by available space) including, by way of
25 example, operating systems, application programs, data files, directory structures, and computer-executable instructions.

[0078] Output devices, such as display device 1102, may be coupled to computing device 1101 via an interface, such as video adapter 1111. Other types of output devices may include printers, audio outputs, tactile devices or other sensory
30 output mechanisms, or the like. Output devices may enable computing device 1101 to interact with human operators or other machines or systems. A user may interface with computing environment 1100 via any number of different input devices 1103 such as a keyboard, mouse, joystick, game pad, data port, and the like. These and other input devices may be coupled to processor 1107 via input/output interfaces

5 1112 which may be coupled to system bus 1108, and may be coupled by other interfaces and bus structures, such as a parallel port, game port, universal serial bus ("USB"), fire wire, infrared port, and the like.

[0079] Computing device 1101 may operate in a networked environment via communications connections to one or more remote computing devices through one
10 or more local area networks ("LAN"), wide area networks ("WAN"), storage area networks ("SAN"), the Internet, radio links, optical links and the like. Computing device 1101 may be coupled to a network via network adapter 1113 or the like, or, alternatively, via a modem, digital subscriber line ("DSL") link, integrated services digital network ("ISDN") link, Internet link, wireless link, or the like.

15 [0080] Communications connection 1114, such as a network connection, typically provides a coupling to communications media, such as a network. Communications media typically provide computer-readable and computer-executable instructions, data structures, files, program modules and other data using a modulated data signal, such as a carrier wave or other transport mechanism.
20 The term "modulated data signal" typically means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communications media may include wired media, such as a wired network or direct-wired connection or the like, and wireless media, such as acoustic, radio frequency, infrared, or other wireless
25 communications mechanisms.

[0081] Those skilled in the art will realize that storage devices utilized to provide computer-readable and computer-executable instructions and data can be distributed over a network. For example, a remote computer or storage device may store computer-readable and computer-executable instructions in the form of
30 software applications and data. A local computer may access the remote computer or storage device via the network and download part or all of a software application or data and may execute any computer-executable instructions. Alternatively, the local computer may download pieces of the software or data as needed, or distributively

5 process the software by executing some of the instructions at the local computer and some at remote computers and/or devices.

[0082] Those skilled in the art will also realize that, by utilizing conventional techniques, all or portions of the software's computer-executable instructions may be carried out by a dedicated electronic circuit such as a digital signal processor
10 ("DSP"), programmable logic array ("PLA"), discrete circuits, and the like. The term "electronic apparatus" may include computing devices or consumer electronic devices comprising any software, firmware or the like, or electronic devices or circuits comprising no software, firmware or the like.

[0083] The term "firmware" typically refers to executable instructions, code
15 or data maintained in an electronic device such as a ROM. The term "software" generally refers to executable instructions, code, data, applications, programs, or the like maintained in or on any form of computer-readable media. The term "computer-readable media" typically refers to system memory, storage devices and their associated media, communications media, and the like.

5

CLAIMS

1. A method for providing an application with fault-tolerant network communications between a plurality of nodes, comprising:

providing a plurality of initial communications pathways via a plurality of networks coupled to the plurality of nodes;

10 receiving a data packet on a sending node from the application, the sending node being one of the plurality of nodes, the data packet being addressed by the application to an address on one of the plurality of nodes; and

selecting a first selected pathway for the data packet from among the plurality of initial communications pathways where the first selected pathway is a
15 preferred pathway.

2. The method of claim 1, further comprising:

detecting a local connectivity event associated with one of the plurality of initial communications pathways; and

indicating whether or not the one of the plurality of initial communications
20 pathways is operational based on the local connectivity event.

3. The method of claim 1, further comprising:

sending a routable heartbeat packet over one of the plurality of initial communications pathways;

monitoring a response to the routable heartbeat packet to determine an end-
25 to-end operational status of the one of the plurality of initial communications pathways; and

indicating whether or not the one of the plurality of initial communications pathways is end-to-end operational based on the monitoring of the response.

30 4. The method of claim 1, wherein the address is an Internet Protocol version 4 address or an Internet Protocol version 6 address.

5. The method of claim 1, wherein the data packet received from the application is a Transmission Control Protocol packet or a User Datagram Protocol packet.

- 5 6. The method of claim 1, wherein computer-executable instructions for performing the method are stored on a computer-readable medium.
7. The method of claim 1, wherein the address is a virtual address.
8. The method of claim 7, further comprising routeably tunneling the data packet over the first selected pathway without requiring the application to be aware of which of the plurality of initial communications pathways is the selected pathway.
- 10 9. The method of claim 8, wherein the data packet routeably tunneled over the first selected pathway is a Transmission Control Protocol packet or a User Datagram Protocol packet.
10. The method of claim 7, further comprising:
- 15 detecting a failure of the first selected pathway;
- selecting a second selected pathway from among the plurality of initial communications pathways where the second selected pathway is both end-to-end operational and a preferred pathway; and
- routeably tunneling the data packet over the second selected pathway
- 20 without requiring the application to be aware of which of the one or more initial communications pathways is the second selected pathway.
11. The method of claim 1, wherein the first selected pathway is selected based on a physical destination address included with the data packet.
12. The method of claim 11, further comprising sending the data packet over the
- 25 first selected pathway without requiring the application to be aware of which of the plurality of initial communications pathways is the first selected pathway.
13. The method of claim 11, further comprising:
- detecting a failure of the first selected pathway;
- selecting a second selected pathway from the plurality of initial
- 30 communications pathways where the second selected pathway is both end-to-end operational and a preferred pathway; and
- routeably tunneling the data packet over the second selected pathway without requiring the application to be aware of which of the one or more initial communications pathways is the second selected pathway.

- 5 14. The method of claim 13, wherein the data packet routeably tunneled over the second selected pathway is a Transmission Control Protocol packet or a User Datagram Protocol packet.
15. A method for providing an application with fault-tolerant network communications for between a plurality of nodes, comprising:
- 10 providing a plurality of initial communications pathways via a plurality of networks coupled to the plurality of nodes;
- receiving a data packet at a receiving node, the receiving node including a fault-tolerant network stack and being one of the plurality of nodes, the data packet being destined to the application;
- 15 determining if the data packet was routeably tunneled; and
- if the data packet was routeably tunneled, forwarding the data packet up the fault-tolerant network stack.
16. The method of claim 15, further comprising:
- receiving a routable heartbeat packet at the receiving node over one of the
- 20 plurality of initial communications pathways; and
- replying with a response to the routable heartbeat packet to indicate an end-to-end operational status of the one of the plurality of initial communications pathways.
17. The method of claim 15, wherein the data packet is a Transmission Control
- 25 Protocol packet or a User Datagram Protocol packet.
18. The method of claim 15, wherein computer-executable instructions for performing the method are stored on a computer-readable medium.

5 19. A system for providing an application with fault-tolerant network communications between a plurality of nodes, comprising:

a first fault-tolerant driver coupled to a first network stack and operating on a first node, the first node being one of the plurality of nodes;

a second fault-tolerant driver coupled to a second network stack and
10 operating on a second node, the second node being one of the plurality of nodes;
and

the first fault-tolerant driver and the second fault-tolerant driver being coupled via a plurality of initial communications pathways over a plurality of networks;

15 20. The system of claim 19, wherein the first fault-tolerant driver comprises:

a processing element coupled to the application via the first network stack;

a routing database coupled to the processing element including:

an entry representing a pathway to the second node including a
20 physical address of the second node, the pathway being one of the plurality of initial communications pathways; and

an indication of an end-to-end operational status of the pathway to the second node;

a protocol adapter coupled to the processing element and coupled to one of
25 the plurality of networks via the first network stack, the one of the plurality of networks being associated with the pathway to the second node; and

a tunnel adapter associated with the pathway to the second node and coupled to the processing element and coupled to the one of the plurality of networks via the first network stack.