


---

# z/OS Communications Server

## System management and monitoring

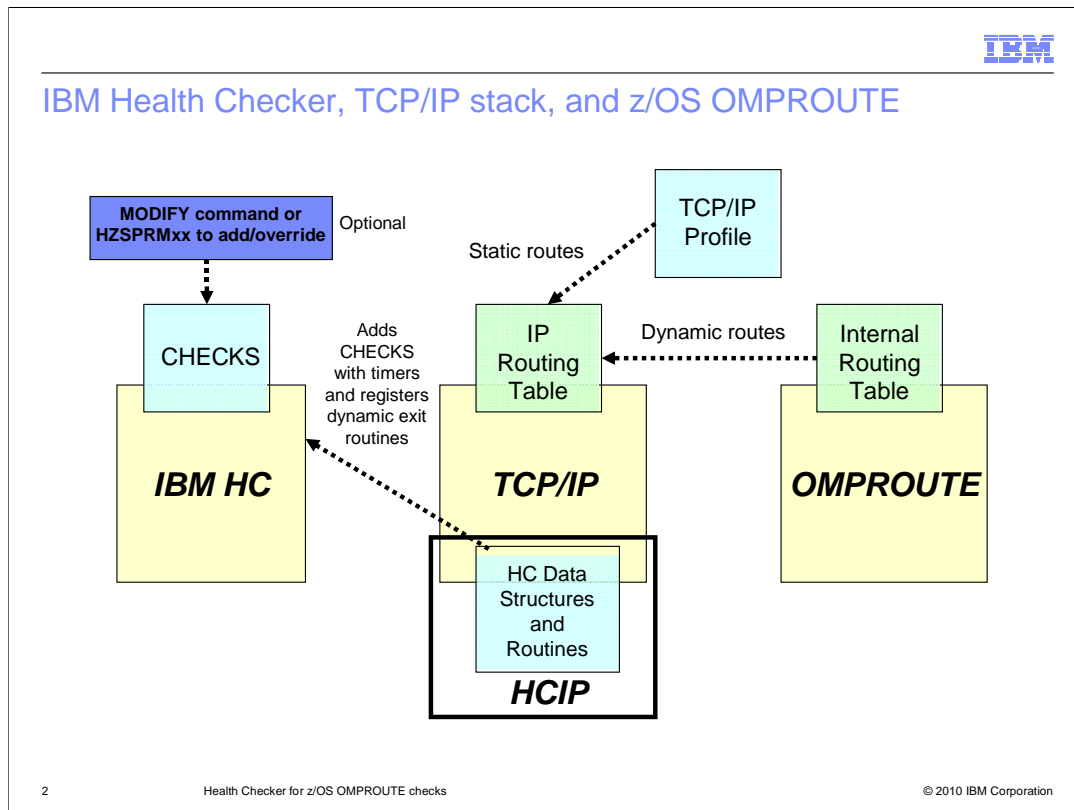
### IBM Health Checker for z/OS OMPROUTE checks



© 2010 IBM Corporation

This presentation describes the new functions in z/OS V1R12 Communications Server that are contained in the system management and monitoring theme for OMPROUTE. Additional health checks have been added to the IBM Health Checker to monitor dynamic routes.

## IBM Health Checker, TCP/IP stack, and z/OS OMROUTE

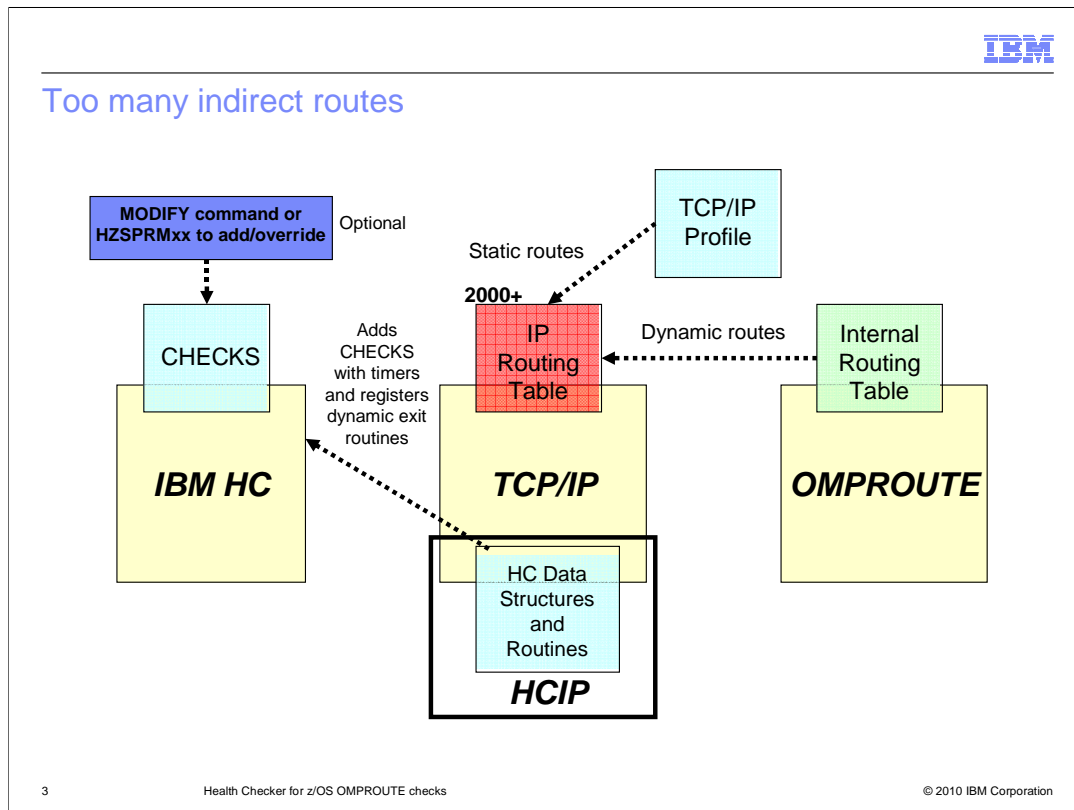


The IBM Health Checker is a z/OS component that helps identify potential configuration and migration compliance problems before they impact availability or cause outages. Individual products, z/OS components, or Independent Software Vendor software can provide checks that take advantage of the IBM Health Checker for z/OS framework.

The TCP/IP stack has IPv4 and IPv6 routing tables which can consist of both static and dynamic routes. OMROUTE and TCP/IP profile statements dynamically update the routing tables. OMROUTE is a multi-protocol routing daemon that provides dynamic routing support and manages a TCP/IP stack's routing table by sending route add and delete requests. It supports RIPv1, RIPv2, RIPv6, OSPFv2, and OSPF for IPv6 protocols. It also supports configured IPv4 and IPv6 static routes defined in the TCP/IP stack.

The slide shows a diagram of interacting components: IBM Health Checker for z/OS (IBMHC), TCP/IP stack, TCP/IP Health Checker (HCIP), and OMROUTE. IBMHC, TCP/IP stack, and OMROUTE run in their own address spaces. During initialization, the TCP/IP stack defines and adds the TCP/IP checks and timer parameters to the IBMHC. Based on these parameters, IBMHC runs the checks for health monitoring. TCP/IP configuration processing updates the IPv4 and IPv6 static routes in the TCP/IP stack routing tables. OMROUTE also updates the IPv4 and IPv6 dynamic routes in the TCP/IP stack routing tables. IBMHC interacts with HCIP to update or retrieve health checker related data. Based on the check timer parameters and health checker data from HCIP, IBMHC issues informational or warning messages as necessary.

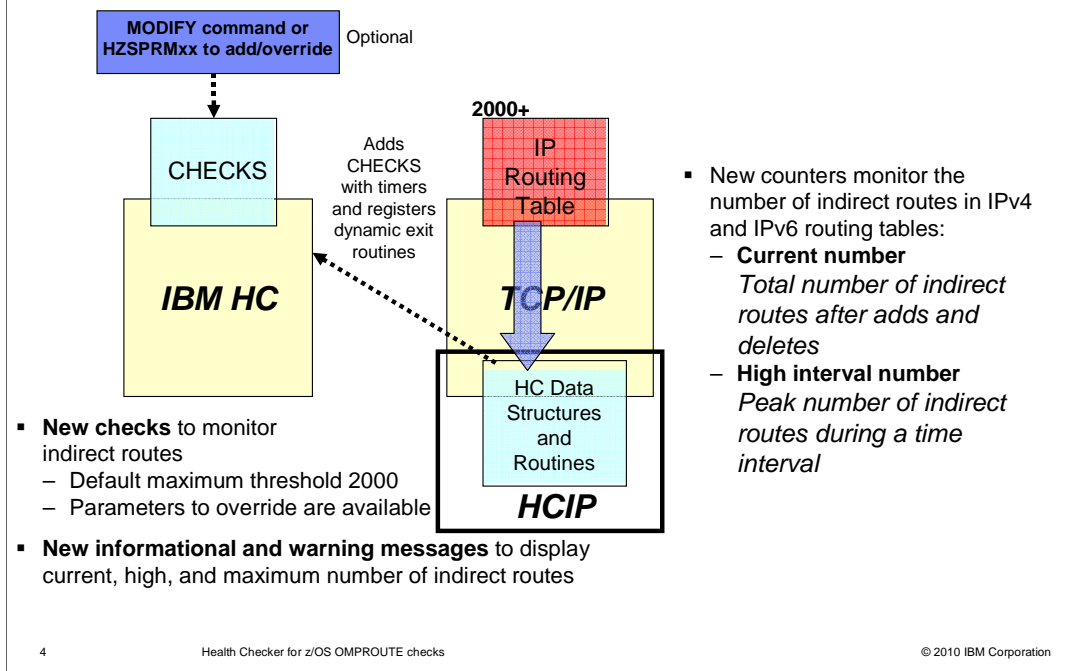
## Too many indirect routes



A routing table that is considered to be excessive (2000 routes or more) can cause inefficiency in network design and less than optimal performance for OMPROUTE and TCP/IP. Most z/OS sites appear to have 50-500 unique routes. IBM service frequently tells customers with more than 2000 routes to reduce the number of routes after determining that performance degradations in OMPROUTE and TCP/IP were caused by the excessive number of routes. There have been a small number of customers over the years who have attempted to configure many thousands of routes (from both dynamic and static routing protocols) on z/OS when they only needed 100 or so. Most of the time, having many thousands of routes does not cause a problem. However, if all of the routes ever need to be deleted or added at the same time, then high processor consumption might be seen in the TCP/IP stack or in OMVS. Many thousands of routing updates have to be processed to make the routing changes.

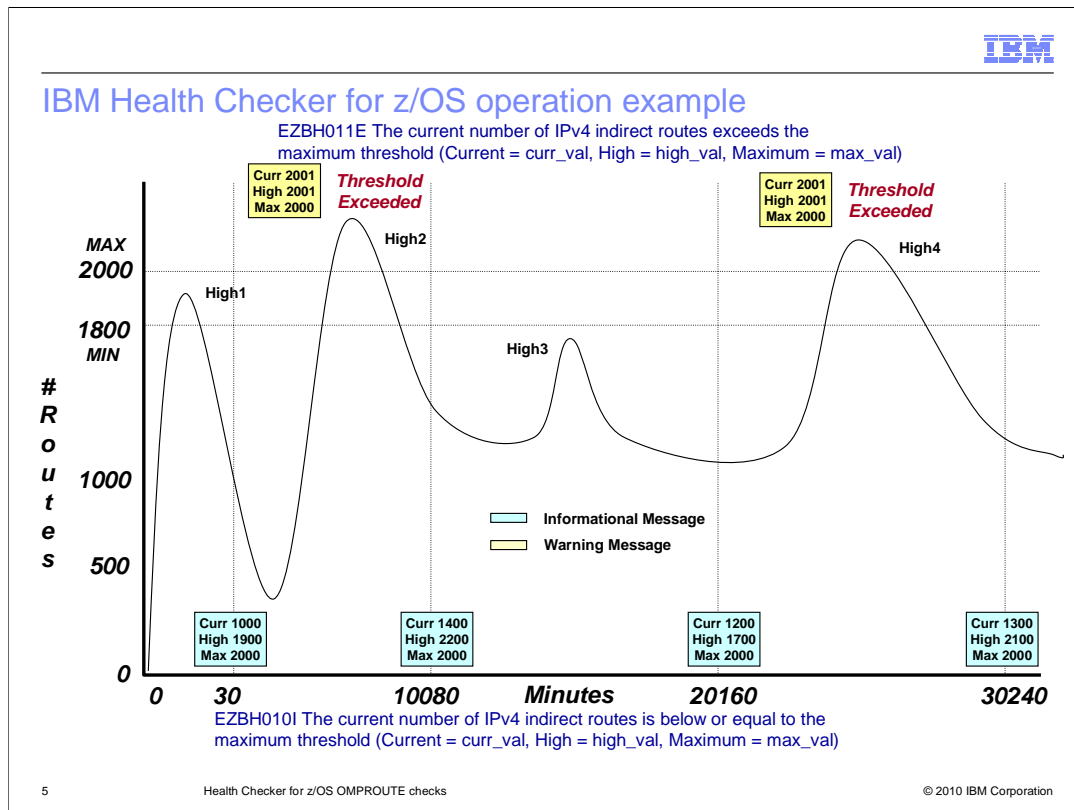
For rapid responsiveness to network topology changes and for fast route convergences, the OSPF protocol uses short-interval timers to maintain adjacencies and to make necessary routing changes. The Hello Interval timer is used for adjacency attempts and to maintain adjacencies with neighbors. The Dead Router Interval timer is used to make routing changes. Because short-interval timers are used, the time to process the many thousands of routing updates might exceed the OSPF dead router intervals. This results in OSPF adjacency losses with neighbors and contributes to network connectivity problems.

## IBM Health Checker for z/OS OMPROUTE checks



New counters have been added to keep track of the total and peak number of indirect routes in the routing table. These counters are used by IBMHC for the health check monitoring and for input into the informational and warning messages. The current number is incremented and decremented when the route table is updated. The high interval number is set to the peak number of indirect routes during a time interval and is reset to the current number for the next interval.

IBMHC has been enhanced to monitor the number of IPv4 and IPv6 indirect routes in a TCP/IP stack. The default maximum threshold for each (2000), new check parameters used for the default overrides, and new messages (informational and warning) are defined to IBMHC.



IBMHC performs the new checks to monitor the number of IPv4 and IPv6 indirect routes. A one-time 30-minute check after TCP/IP initialization and interval checks based on the health checker's INTERVAL setting. The interval timer starts at TCP/IP initialization and no informational or warning messages are issued at this time because updates to the routing tables might not be complete. If a counter exceeds the maximum threshold value, HCIP notifies the IBMHC which immediately issues a warning message. If a maximum threshold value was dynamically modified by an operator, IBMHC immediately issues a warning or informational message to report the change. If IBMHC has detected no exceptions, it issues informational messages. If a message threshold has been exceeded within a time interval, IBMHC issues a warning message and suppresses subsequent warning messages in that time interval to prevent message thrashing. However, if the number of indirect routes has dropped to below 90% of the maximum threshold in that time interval, warning messages are again issued if the maximum threshold has been exceeded.

This example illustrates informational data (in blue) and warning data (in yellow). The maximum threshold is 2000 and the minimum threshold is 1800 (90% of 2000). The messages display the current, high, and maximum values at the intervals. The warning messages are issued when the number of indirect routes exceeds the maximum threshold. The high interval numbers (High1 = 1900, High2 = 2200, High3 = 1700, and High4 = 2100) are recorded as peak numbers and are reported at the time intervals to provide historical data information.

The informational message, EZBH010I, is shown at the bottom of the slide. This message is not written to the system console because of its successful status. The warning message, EZBH011E, is shown at the top of the slide. This message is written to the system console because of its exception status. Either message type can be directed to a stream data set for a history log. The most recent message of either type can be displayed in SDSF with the CK option by selecting the check in question. Included are an explanation and actions to take for resolution.

## IBM Health Checker for z/OS OMPROUTE considerations

- **If you have a large number of indirect routes:**
  - Disable the checks or increase the maximum thresholds for environment
  - When applicable, use OSPF stub areas, route summarization, or filters to reduce the number of indirect dynamic routes
- **Route growth and drop-off might create exception messages:**
  - OBEYFILE for BEGINROUTES or GATEWAY statement with large number of indirect static routes for full route table replacements
  - Interface or route flapping that influence changes to indirect dynamic routes

If more than 2000 routes are installed on a system and IBMHC is enabled for that system environment, exception messages might be displayed at the system console. To suppress the exception messages, update the POLICY statements in the HZSPRMxx parmlib member and either have the new checks disabled by specifying INACTIVE status or increase the maximum threshold values by overriding the installation defaults.

There are conditions where exception messages might appear suddenly for brief periods of time. Successive route growths followed by drop-offs might be indications of large route table updates. For example, OBEYFILES for BEGINROUTE or GATEWAY statements are issued for full route table replacements. Another possible cause is interface or route flapping in a z/OS system or network such that it influences changes to the dynamic indirect routes.

Interface flapping can occur when the interface state changes as result of some hardware failure on the z/OS system or in the network.

Route flapping occurs when a router alternately advertises a destination network in one route and then another in quick sequence or appear unavailable and then available again.

Interface and route flapping can influence the number of indirect routes from the route deletes and adds.



## IBM Health Checker for z/OS OMPROUTE considerations, continued

- **The one-time 30-minute check does not occur if:**
  - IBMHC is started at least 30 minutes after TCP/IP initialization
  - The specified INTERVAL is less than the 30-minute check
  
- **When using SDSF option CK:**
  - Most recent informational or warning message is displayed
  - Warning messages are also written to system console
  - Messages can be directed to a log stream data set for history log

The one-time 30-minute check depends on the enablement of IBMHC, the time intervals for the new checks, and whether previous warning messages have been issued. There are many conditions where the 30-minute check is not done.

IBMHC provides a utility (SDSF option CK) to display all defined health checks. This option displays only the most recent message. Note that IBMHC also writes warning messages to a system console for alerts. To avoid flooding the system console, informational messages are not written to the system console. However for historical recording purposes, IBMHC can be configured to direct the messages to a log stream data set.

## Policy statements in HZSPRMxx example

- **Example policy statements in SYS1.PARMLIB(HZSPRMxx)**
  - Check names and parameters to override installation defaults for all TCP/IP stacks

### ADDREPLACE POLICY STMT(IPMAXRT4)

```
UPDATE
CHECK(IBMCS,CSTCP_IPMAXRT4_*)
REASON('Your reason')
DATE(yyymmdd)
PARM('IPMAXRT4(1000)')
ACTIVE
SEVERITY(LOW)
INTERVAL(24:00)
```

### ADDREPLACE POLICY STMT(IPMAXRT6)

```
UPDATE
CHECK(IBMCS,CSTCP_IPMAXRT6_*)
REASON('Your reason')
DATE(yyymmdd)
PARM('IPMAXRT6(1000)')
ACTIVE
SEVERITY(LOW)
INTERVAL(24:00)
```

The TCP/IP stack defines the new checks CSTCP\_IPMAXRT4\_tcpipstackname and CSTCP\_IPMAXRT6\_tcpipstackname along with their parameters during TCP/IP initialization. The TCP/IP stack sets up the dynamic exits and loads the checks and parameters into IBMHC to enable the monitoring of those checks. If IBMHC is not active during TCP/IP initialization, the load attempts by the TCP/IP stack are performed again once IBMHC becomes available. When a TCP/IP stack terminates, the checks associated with that TCP/IP stack are removed from IBMHC.

New parameters IPMAXRT4(n) and INTERVAL(hhh:mm) for check CSTCP\_IPMAXRT4\_tcpipstackname and new parameters IPMAXRT6(n) and INTERVAL(hhh:mm) for check CSTCP\_IPMAXRT6\_tcpipstackname are added by HCIP to provide the defaults to IBMHC. For a maximum threshold value, n ranges from one to 65536 and defaults to 2000. For a time interval, hhh ranges from 0 to 999 hours and mm ranges from 0 to 59 minutes. The default time interval is 168 hours and 0 minutes (168:00) which is seven days. These parameters can also be used in POLICY statements and MODIFY commands to override the default values.

The POLICY statements in a HZSPRMxx parmlib member can be used to override the installation defaults provided by the TCP/IP stacks. In this example, the wildcard suffixes in the check names, denoted by asterisk (\*), indicate that the checks and their parameters apply to all TCP/IP stack job names on this system. In addition, the maximum thresholds of 1000 and time intervals of 24 hours for the checks are specified to override the defaults (maximum thresholds of 2000 and time intervals of 168 hours). The ACTIVE parameters (instead of INACTIVE) are specified to indicate that the checks are to be enabled.



## MODIFY HZSPROC examples

- **Example MODIFY commands to override current or default values**

```
F hzsproc,update,check=(ibmcs,cstcp_ipmaxrt4_*),parm='ipmaxrt4(1500)'
F hzsproc,update,check=(ibmcs,cstcp_ipmaxrt6_*),interval=01:30
F hzsproc,update,check=(ibmcs,cstcp_ipmaxrt4_*),interval=00:30,parm='ipmaxrt4(1000)'
F hzsproc,update,check=(ibmcs,cstcp_ipmaxrt6_*),interval=00:30,parm='ipmaxrt6(1000)'
```

- **Example MODIFY command to display a check summary report**

```
F hzsproc,display,check=(ibmcs,cstcp_ipmaxrt4_*)
HZS0200I 20.43.34 CHECK SUMMARY 336
CHECK OWNER CHECK NAME STATE STATUS
IBMCS CSTCP_IPMAXRT4_TCPCS2 AE SUCCESSFUL
IBMCS CSTCP_IPMAXRT4_TCPCS AE SUCCESSFUL
A - ACTIVE I - INACTIVE
E - ENABLED D - DISABLED
G - GLOBAL CHECK + - CHECK ERROR MESSAGES ISSUED
```

The MODIFY HZSPROC command is used to update the check parameters and the check time intervals for the new CSTCP\_IPMAXRT4\_\* and CSTCP\_IPMAXRT6\_\* checks. A MODIFY command can be used to update a check for all TCP/IP stacks (using the wildcard, asterisk) or for a particular TCP/IP stack by name (for example, TCPCS). The MODIFY commands can be used to override the current or default values. Note that the alternative method to a MODIFY command is to dynamically update the corresponding field for a check in the SDSF option CK display.

The MODIFY command displays the CHECK SUMMARY report for the new check CSTCP\_IPMAXRT4\_\* for all TCP/IP stacks. The report lists the check's owner, name, state, and status. The MODIFY command syntax is similar for the CSTCP\_IPMAXRT6\_\* check. A TCP/IP stack name can be specified to display a CHECK SUMMARY report for that TCP/IP stack. In the example for a successful MODIFY command for the CSTCP\_IPMAXRT4\_\* check, note that IBMHC issues an informational message at the operator console to indicate that the update was successful.

Another MODIFY command, not shown, displays the CHECK DETAIL report for the new check CSTCP\_IPMAXRT4\_\* for all TCP/IP stacks. Basically, the report provides more detailed information about a check such as the configuration information on INTERVAL, DEFAULT PARAMETERS, and USER SUPPLIED PARAMETERS. The MODIFY command syntax is similar for the CSTCP\_IPMAXRT6\_\* check. A TCP/IP stack name can be specified to display a CHECK DETAIL report for that TCP/IP stack.



## SDSF option CK (S.CK) display example

```

Display Filter View Print Options Search Help
-----
SDSF HEALTH CHECKER DISPLAY MVS052                LINE 21-56 (135)
COMMAND INPUT ==>                                SCROLL ==> CSR
NP  NAME                                           State           Status
    CSTCP_IPMAXRT4_TCPCS                          ACTIVE (ENABLED) SUCCESSFUL
    CSTCP_IPMAXRT6_TCPCS                          ACTIVE (ENABLED) SUCCESSFUL
    ...
F1=HELP      F2=SPLIT      F3=END      F4=RETURN      F5=IFIND      F6=BOOK
F7=UP        F8=DOWN        F9=SWAP     F10=LEFT       F11=RIGHT     F12=RETRIEVE

```

- PF11 once - Start time, interval, and next scheduled time and date are shown

```

NP  NAME                                           Start-Date-Time Interval NextSch-Date
    CSTCP_IPMAXRT4_TCPCS                          01/04/2010 20:43:20 168:00 01/04/2010 20:
    CSTCP_IPMAXRT6_TCPCS                          01/04/2010 20:43:20 168:00 01/04/2010 20:

```

- PF11 twice – Check parameters

```

NP  NAME                                           Len Parameters           SysLev
    CSTCP_IPMAXRT4_TCPCS                          14 ('IPMAXRT4(2000)')   z/OS 0
    CSTCP_IPMAXRT6_TCPCS                          14 ('IPMAXRT6(2000)')   z/OS 0

```

10

Health Checker for z/OS OMPROUTE checks

© 2010 IBM Corporation

This slide has a screen capture of an SDSF Health Checker Display for the list of checks. The display shows the new checks CSTCP\_IPMAXRT4\_TCPCS and CSTCP\_IPMAXRT6\_TCPCS. IBMHC has determined that these checks have been successful. The active check states are also displayed here.

Press the PF11 key to scroll to the right. Here, a check name field is shown with the check's start date and time, time interval, and the next scheduled date and time. Note that the value in the time interval field can be changed dynamically by overtyping it with a new one.

Press the PF11 key again to scroll farther to the right to see more parameters.

## SDSF option CK (S.K) display – Select (continued)

- If a status is exception, you see a warning message:

```

Display Filter View Print Options Search Help
-----
SDSF OUTPUT DISPLAY CSTCP_IPMAXRT4_TCPCS          LINE 0          COLUMNS 02- 81
COMMAND INPUT ==>                                SCROLL ==> CSR
***** TOP OF DATA *****
CHECK(IBMCS,CSTCP_IPMAXRT4_TCPCS)
START TIME: 01/04/2010 21:23:07.592165
CHECK DATE: 20100104 CHECK SEVERITY: LOW
CHECK PARM: IPMAXRT(2000)

* Low Severity Exception *

EZBH011E The current number of IPv4 indirect routes exceeds the maximum
threshold (Current = 2001, High = 2001, Maximum = 2000)

Explanation: The CSTCP_IPMAXRT4_TCPIPstackname check determined that ...
System Action: The system continues processing.
Operator Response: Contact the system programmer.
System Programmer Response: To reduce the number of indirect OSPF routes ...
Problem Determination: Not applicable.
Source: z/OS Communications Server Health Checker
...
END TIME: 01/04/2010 21:23:07.598584 STATUS: EXCEPTION-LOW
***** BOTTOM OF DATA *****

```

Select a check by typing the letter ‘S’ next to the check’s name. A new panel opens to display any informational or warning message.

The slide has a screen capture of a new SDSF panel displaying the warning message for the selected check CSTCP\_IPMAXRT4\_TCPCS with an EXCEPTION-LOW status. This indicates that the IBM Health Checker for z/OS has determined that the maximum threshold for the check has been exceeded. Note that the warning message displays the current, high interval, and maximum threshold values relating to the number of indirect routes. Also, the display includes the message explanation and details on actions to take for resolution. For a CSTCP\_IPMAXRT6\_\* check with an EXCEPTION-LOW status, the panel for the EZBH013E message has a similar appearance. A similar panel is shown for informational messages.



## Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback about SysMgmtOMPROUTE.ppt](mailto:iea@us.ibm.com?subject=Feedback%20about%20SysMgmtOMPROUTE.ppt)

This module is also available in PDF format at: [../SysMgmtOMPROUTE.pdf](http://../SysMgmtOMPROUTE.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, Current, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2010. All rights reserved.