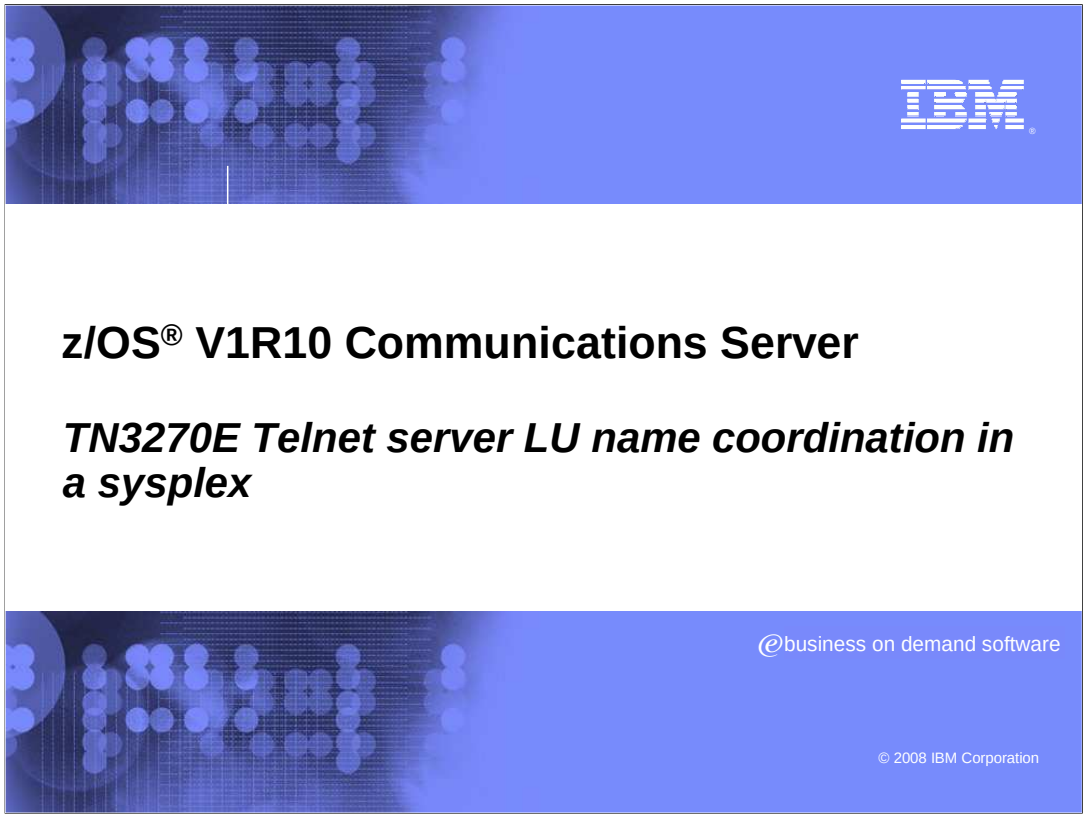

This presentation discusses TN3270E Telnet server LU name coordination in a sysplex.

TN3270E Telnet server LU name coordination in a sysplex background

- Telnet uses an LU to represent each client emulator
  - Each LU name must be unique
  - A duplicate LU name will fail the open ACB request in VTAM®

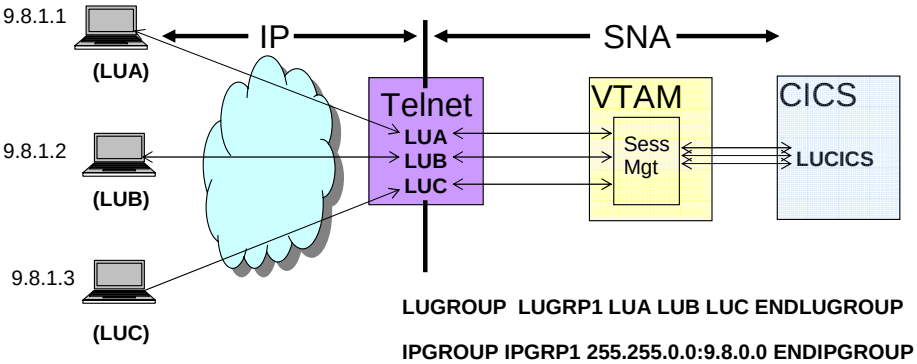

LUGROUP LUGRP1 LUA LUB LUC ENDLUGROUP

IPGROUP IPGRP1 255.255.0.0:9.8.0.0 ENDIPGROUP

LUMAP LUGRP1 IPGRP1

The TN3270E Telnet server, Telnet, is a gateway between the IP network and the SNA network. Telnet accepts IP connections from client emulators and establishes an SNA session through VTAM to extend the data path to the target application. An SNA session consists of the target application, which is a primary Logical Unit (LU), and a terminal or another application, which is a secondary LU. Telnet creates a unique single-session application secondary LU for its half of the SNA session to represent each client connection. VTAM does not allow duplicate LU names. If Telnet tries to create a duplicate LU name during the OPEN ACB process, VTAM fails the creation attempt.

In the figure on this slide, Telnet has defined a group of LUs which is named LUGROUP1. A group of IP addresses, which is any address that falls into the 9.8.0.0/16 subnet is also defined and named IPGROUP1. Then an LUMAP statement tells telnet to use LUs from LUGROUP1 to represent IP addresses from IPGROUP1. So when a client on 9.8.1.1 requests a session with CICS, Telnet assigns LUA to represent that client. LUB is assigned when the client on 9.8.1.2 comes in, and LUC is assigned to the client on 9.8.1.3. Once an LU name is assigned to a client, that name cannot be reused because that would result in duplicate LU names on the SNA side of the session and cause the OPEN ACB to fail.

Duplicate LU name assignments

- The same LU name can be concurrently assigned to different Telnet clients by different Telnet servers

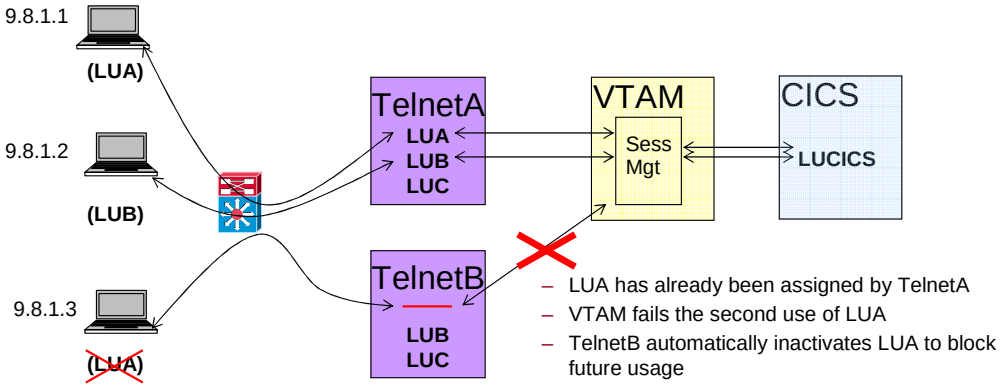

A single-session application LU name can only be used for one SNA session at a time. A single Telnet server has logic to prevent the duplicate creation of an LU name for more than one active client connection. However, V1R9 and earlier Telnet servers are independent and unaware of LU name assignments made by other Telnet servers.

The diagram on this slide illustrates the problem with using sysplex distributor to load balance Telnet clients across several Telnet servers with identical profiles. Each Telnet server maps unique LU names to its own clients, but it is unaware that another Telnet server has also mapped the same LU name to one of its clients. When the server attempts to initiate an SNA session for the client, some of the OPEN ACBs fail with LU name in use. The Telnet server marks that LU name as INACTIVE in its LU name tree so it won' t be used again.

V1R9 duplicate LU name management

- In V1R9 and earlier releases, manual partitioning of the LU name list is required when running multiple Telnet servers

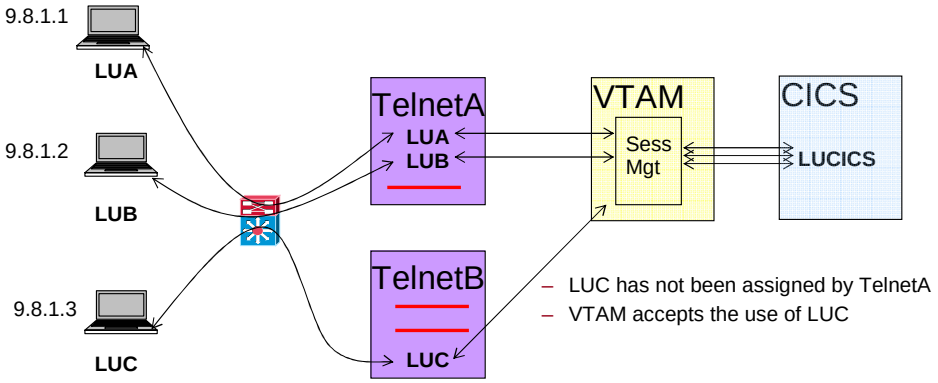

The system administrator must ensure that LU names used at each server can not also be used at another server. This can lead to significant administrative work when configuration updates are needed. In this example, Telnet A can assign LUA and LUB, but not LUC. TelnetB can assign LUC but not LUA or LUB. So in this case the LU namespace is partitioned by manual configuration.

TN3270E Telnet server
LU name coordination in a sysplex

- One Telnet allocates LUs to all other Telnets in a sysplex
  - LUNS – LU Name Server
  - LUNR – LU Name Requester

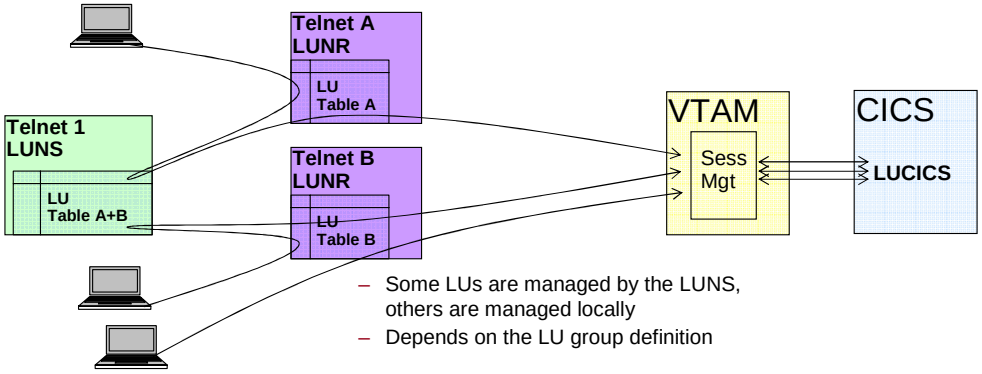

- Some LUs are managed by the LUNS, others are managed locally
- Depends on the LU group definition

V1R10 introduces sysplex-wide coordination of TN3270E LU names. This solution eliminates the need to manually partition the Telnet LU namespace in a sysplex environment. Instead, several telnet servers in a sysplex can now share a namespace, aided by one of them doing coordination and allocating LU names to the other telnets. However the servers are not required to share all of their LU names, some can be kept local to a specific telnet server. When defining LU name groups, a group can be either shared or local.

The Telnet server used for centralized LU allocation is referred to as an LU Name Server, or LUNS (rhymes with tons). Telnet servers that accept client connections will request an LU name from the LUNS instead of directly assigning an LU name. This Telnet is referred to as an LU Name Requester, or LUNR (pronounced lunar). Centralizing LU name allocation ensures no duplicate LU name assignments. All LU groups are defined at the LUNR. Existing LU group definitions are considered local groups. A new set of LU group definitions define shared LU groups. The shared LU group definitions are sent to the LUNS. A Telnet server can have a mixture of shared and local LU groups. Only connections mapped to shared LU groups will use the new LUNS service for LU assignment. If a connection maps to a local LU group, Telnet will assign the LU directly, as it always has.

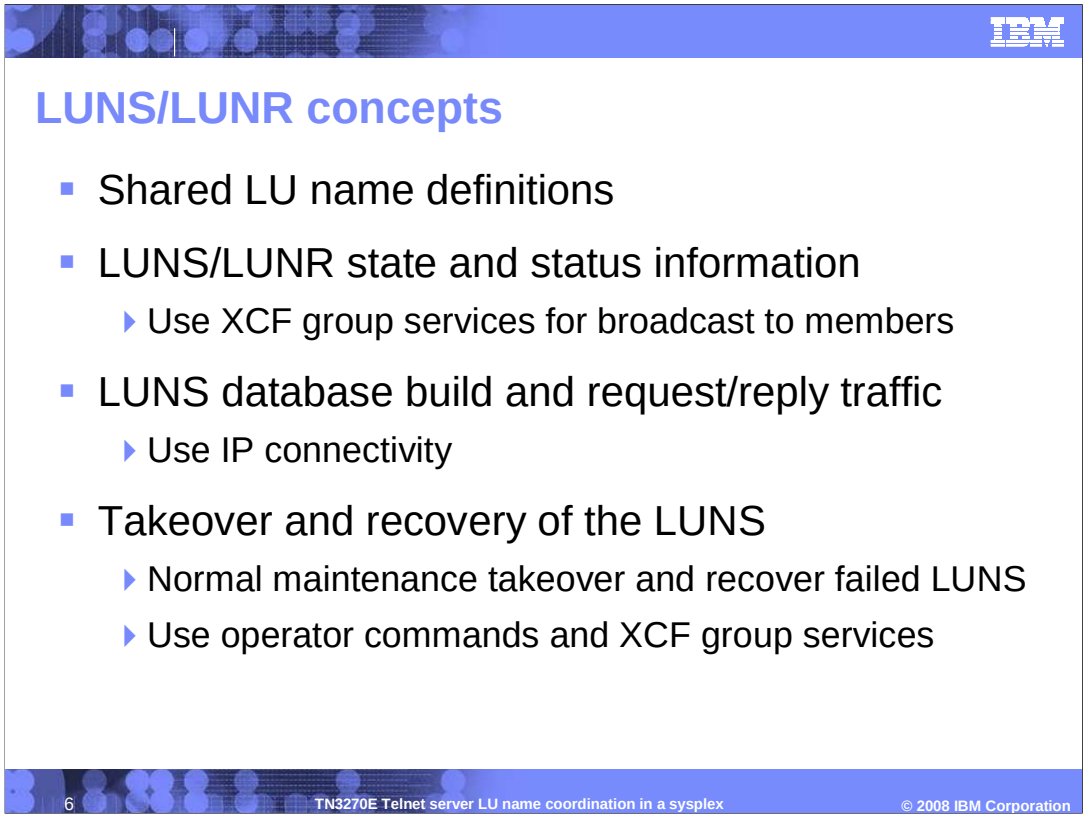

There are several design concepts to briefly consider.

Shared LU groups need to be identified at the LUNR so their definitions can be sent to the LUNS. New shared LU group definitions have been created allowing Telnet to support assigning LU names from local LU groups and assigning LU names allocated from shared LU groups.

The LU Name Server and LU Name Requesters must be aware of the state and status of each other. For example, if a LUNR is activated before a LUNS, the LUNR must know when the LUNS becomes active. The LUNRs must know the IP address and port of the LUNS administrative listener. State and status information is a small amount of data. XCF communication services works well for broadcasting updates to all members. By limiting the information passed over XCF, Telnet needs to use only the XCF group services.

The LUNR owns all LU definitions and must send shared LU group definitions to the LUNS. A potentially large volume of individual LU name requests and replies will flow between the LUNS and LUNR. IP communication handles large amounts of data well.

Takeover of an active LUNS for maintenance and automatic recovery of a failed LUNS are essential requirements of the LUNS/LUNR design. When a new LUNS activates, all the LUNRs must recognize the new LUNS, update the new LUNS with their shared LU group definitions and their currently assigned shared LUs.

Shared LU group definitions

- Shared LU groups are defined and mapped at each LU Name Requester (LUNR)
  - The LU group definitions are sent to the LU Name Server (LUNS) when the profile is created
  - LU names in shared terminal and printer groups are centrally managed by the LUNS within an XCF group
- LU names in non-shared terminal and printer groups remain locally managed at the Telnet server

New LU group object types have been created which are used at the LUNR to indicate that the set of LU names in that group need to be centrally coordinated by the LUNS. These object types are called shared LU groups. Shared LU group definitions can be the same on multiple Telnet servers, but they do not have to be. The shared LU group definitions are sent to the LUNS for central management of LU names. If there is no active LUNS, the LUNR waits and the profile remains in PENDING state until a LUNS becomes active or another obeyfile is processed which purges the current PENDING profile.

Existing non-shared LU groups continue to be supported. Assignment of LU names from non-shared LU groups is managed locally on the Telnet server where the LU group is defined. Other Telnet servers in the sysplex remain unaware of non-shared LU name assignments. LU names in non-shared LU groups still must be manually administered to prevent duplicate client assignment. LU names should not be defined in both shared and non-shared LU name groups. Otherwise, duplicate LU assignments are still possible between non-shared and shared LU name groups.

LUNS management of LU groups

- LUNS allocates LU names to a LUNR only if the LU name is defined **at that LUNR**.

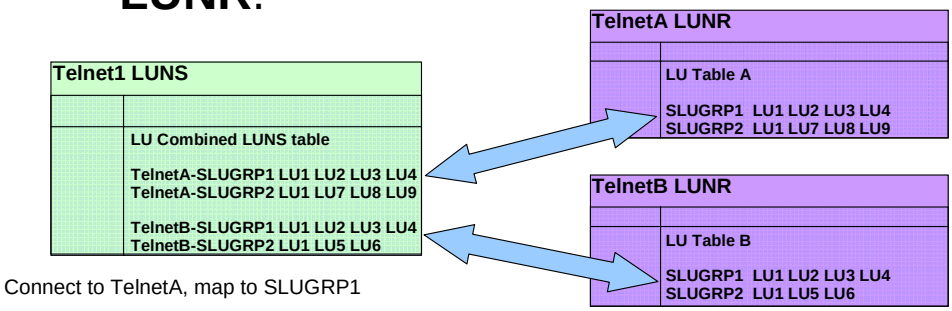

Connect to TelnetA, map to SLUGRP1

LUNS assigns LU1 (LU1 unavailable in all groups)

Connect to TelnetB, map to SLUGRP2

LUNS assigns LU5 (LU5 next available in TelnetB-SLUGRP2)

The Telnet LU Name Server (LUNS) keeps track of LU group definitions by Telnet LU Name Requester (LUNR). The LU name must be defined by a LUNR for the LUNS to allocate the LU name to that LUNR. The LUNR LU name request contains the LU group name that should be searched at the LUNS. If multiple LUNRs define the same LU name or the same LU name is defined in several LU groups, the LUNS ensures only one LUNR is using the LU name.

For example, consider the configuration shown on this slide. If a client connection is accepted at TelnetA and maps to SLUGRP1 the LU name request to the LUNS includes the LUNR name, TelnetA, and the LU group name, SLUGRP1. The LUNS searches only TelnetA-SLUGRP1 for a possible LU allocation and ensures that LU has not been allocated to any other connection across all LUNRs. Assume LU1 is allocated and another client connection is accepted at TelnetB and maps to SLUGRP2. The LU name request to the LUNS will include TelnetB and SLUGRP2 and searches only TelnetB-SLUGRP2 for an available LU name. LU1 is already in use and LU5 is available. The LUNS will allocate LU5 for TelnetB to use.

Shared LU Group statements

- Each of the LU mapping statements now has a shared equivalent:
  - SLUGROUP … ENDSLUGROUP
  - SPRTGROUP … ENDSPRTGROUP
  - SDEFAULTLUS … ENDSDEFAULTLUS
  - SDEFAULTPRT … ENDSDEFAULTPRT
  - SDEFAULTLUSSPEC … ENDSDEFAULTLUSSPEC
  - SDEFAULTPRTSPEC … ENDSDEFAULTPRTSPEC
- The Telnet server must JOIN an XCF group to use shared LU mapping statements.

Before shared LU name support, six statements were available for defining LU groups that can map to client identifiers. Those statements remain and are used to define LU groups that are managed locally by the Telnet server they are defined on.

Each of these local LU group statements now has an equivalent statement used to define shared LU groups. Telnet must join the XCF group to use the shared LU group statements. When shared LU groups are defined, Telnet is considered a LUNR.

Telnet XCF group participation

- TN3270E Telnet server roles (1-2 of 3):
  - Classic Telnet
    - Not a member of a Telnet XCF group
    - No shared LU groups permitted
  - XCF Telnet
    - Must Join XCF Group
    - LUNR capable (Ports can use shared LU groups)

TN3270E Telnet servers can play different roles in a sysplex. If no changes are made to a Telnet server's configuration, it does not join an XCF group and does not participate in coordinated LU name assignment. These servers are now called "Classic" Telnet servers in the documentation.

A TN3270E Telnet server that joins an XCF group can display the status of all of the members in that group and is called an "XCF Telnet" in the documentation. An XCF Telnet is capable of defining shared LU groups for use on a Telnet port. When shared LU groups are defined, the server is called an LU Name Requester or "LUNR" Telnet in the documentation.

Telnet XCF group participation

- TN3270E Telnet server roles (3 of 3)
  - XCF LUNS Telnet
    - Must Join XCF Group
    - LUNR capable
    - Coordinates shared LU name assignments
    - One LUNS in control at any given time
      - others are standby
    - Manual or automated takeover by standby LUNS

A TN3270E Telnet server that joins an XCF group and is configured to coordinate LU name assignments is called an LU Name Server or “LUNS” Telnet in the documentation. An XCF group can optionally have multiple members that are capable of providing the LUNS role to the group. Planned takeover can be initiated manually through a console command. Support for automated takeover in the advent of a LUNS failure is provided.

Any given Telnet server can be configured to participate in any combination of these roles. A LUNR Telnet might have ports that use only shared LU name groups, or a mixture of shared and non-shared LU name groups. A LUNS Telnet can also be a LUNR Telnet or it can be a dedicated LUNS Telnet.

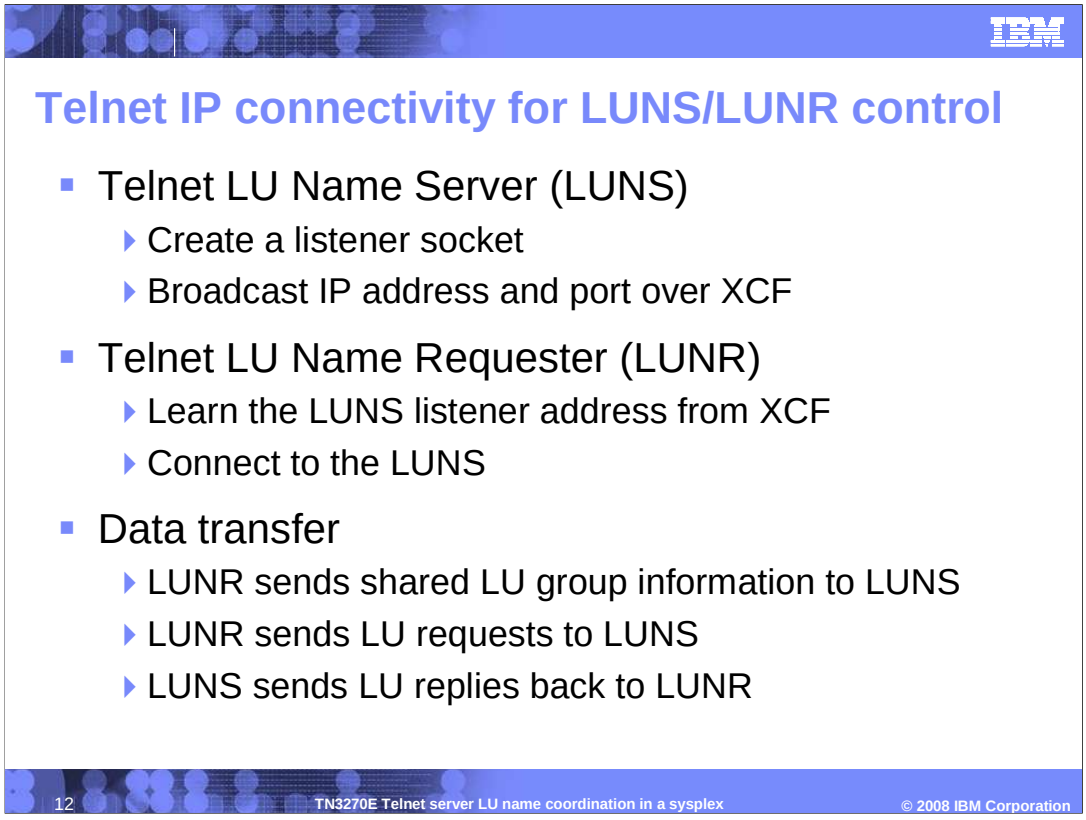

As stated earlier, IP connectivity is used by the LUNS and LUNRs to send and receive LU related data.

The LUNS creates a listener socket for LUNS and LUNR LU administration.

The IP address and port are broadcast by the LUNS over XCF. The LUNR receives the XCF group status update containing the IP address and port. The LUNR connects to the LUNS and the administrative IP connection is established.

Once the IP connection is established, the LUNR sends shared LU group information from its profile to the LUNS. The LUNS uses the LU group information during LU allocation. When client connections are accepted by the LUNR and the connection maps to a shared LU group, the LUNR sends an LU request to the LUNS. The LUNS uses the LU group information it has, determines an available LU name from the central LU table, allocates the LU name to that LUNR, and sends the name to the LUNR. The LUNR processes the name through its own LU table to ensure the name is not already in use locally. If not, the LU is assigned to the connection and normal processing continues.

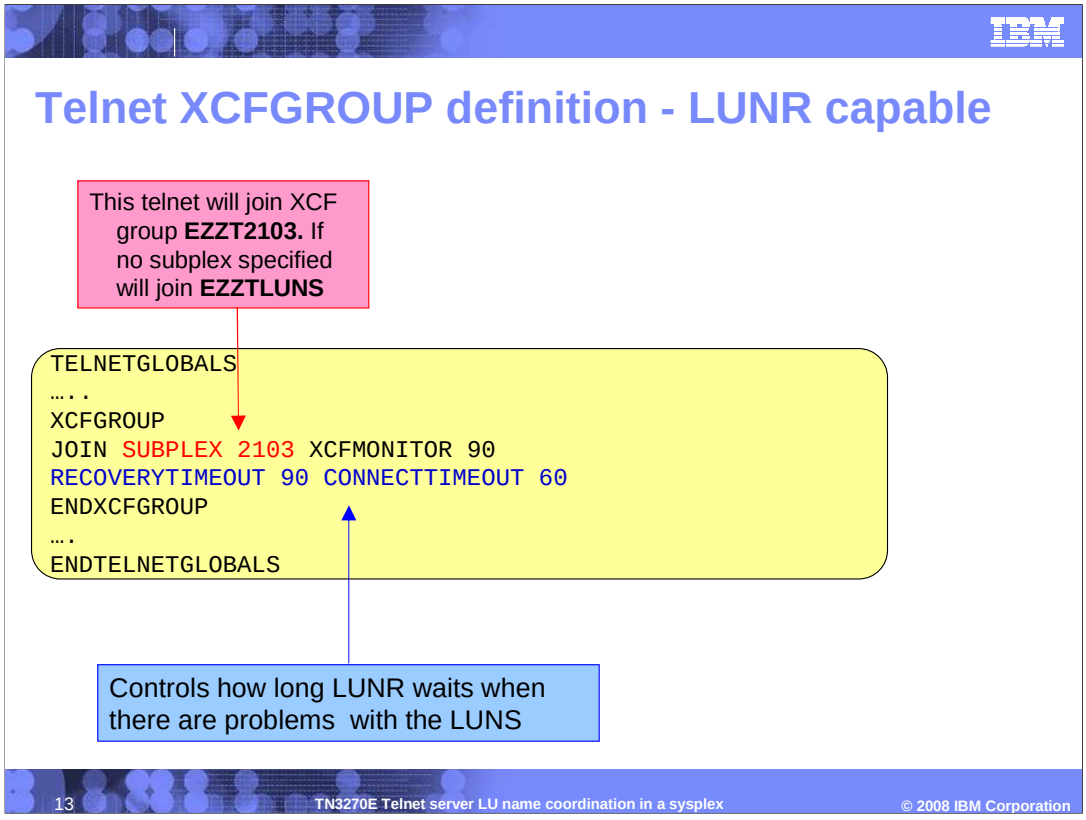

The new XCFGROUP configuration statement determines what roles the TN3270E Telnet server is capable of playing and configures new parameters.

JOIN or **NOJOIN** controls whether the server should join an XCF group. A server must join an XCF group to define and use shared LU name group objects.

The **SUBPLEX suffix** can be used to modify the name of the XCF group to join. The default group name is **EZZTLUNS**. A suffix from one to four characters can be specified. The suffix is right justified and overlays the end of the default name.

XCFMONITOR sets the interval of how often a common storage field is checked by XCF. A Telnet routine updates the field twice every interval period. This routine also examines several internal tasks and the LUNS/LUNR TCP connections. It sets the X flag under the PDMON heading in the XCF GROUP display if a problem is detected.

LUNR-specific parameters control when and what action should be taken if a connection to the LUNS can not be made.

RECOVERYTIMEOUT sets the recovery timeout interval to a specified number of seconds. The default value is 60 seconds. It can be turned off by specifying 0. See the “Problem Detection Monitor” slide for usage details.

CONNECTTIMEOUT sets the connect timeout interval to a specified number of seconds. The default value is 60 seconds. It can be turned off by specifying 0. See the “Problem Detection Monitor” slide for usage details.

See next slide for an example of a LUNS-capable telnet definition.

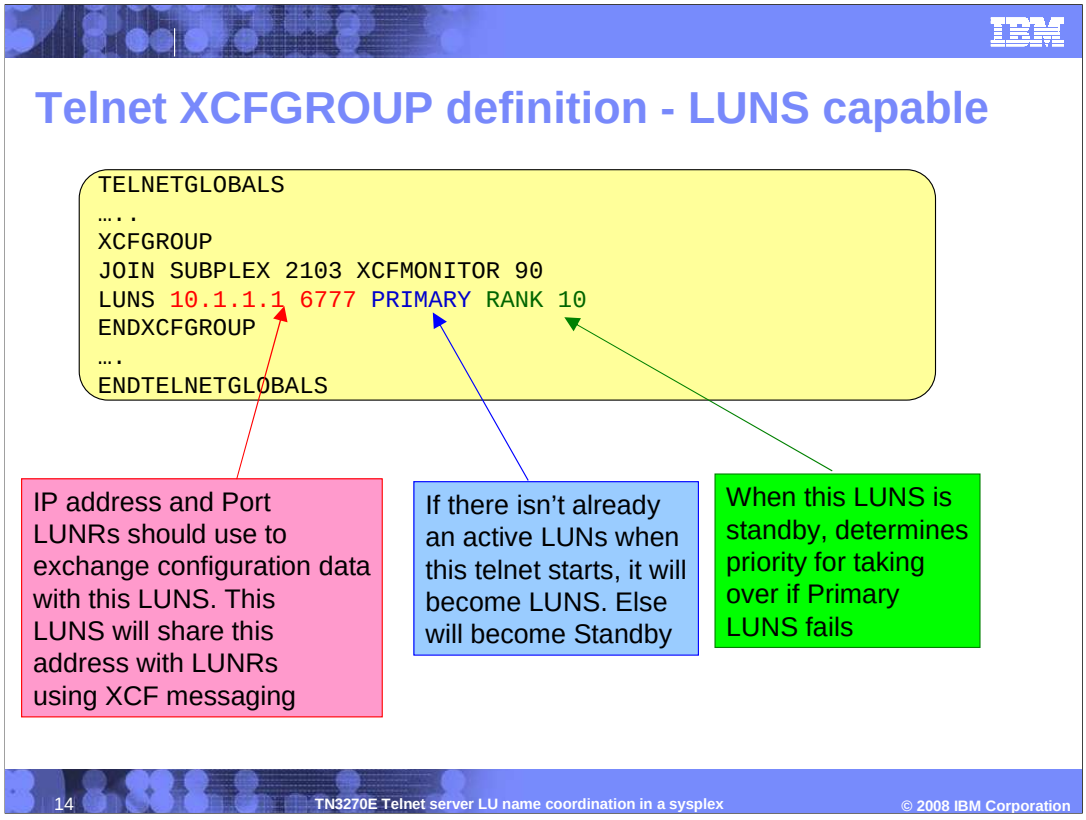

This slide shows an XCFGROUP definition for a LUNS capable Telnet.

Telnet must join the XCF group and specify LUNS-specific parameters to be a LUNS.

The ***ipaddr*** and ***port*** parameters specify where this Telnet will listen for LUNR connection requests when it becomes the active LUNS.

PRIMARY specifies that this Telnet will become the active LUNS at job initiation if there is not already an active LUNS.

BACKUP specifies that this Telnet will be in LUNS STANDBY state at job initiation.

RANK specifies the RANK used at recovery time when an active LUNS fails. The standby LUNS with the highest rank will become the new LUNS.

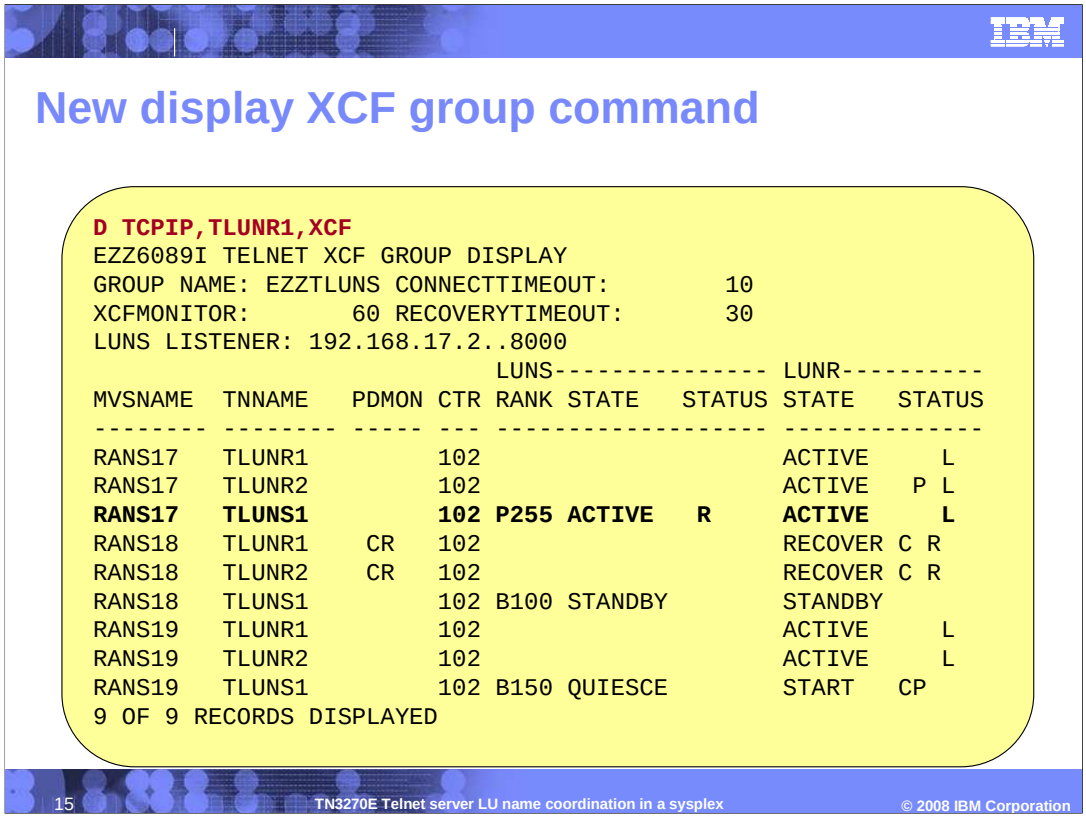

A new display option, XCF, is provided on the DISPLAY TELNET command to help you monitor and debug the telnet XCF and LUNS/LUNR state.

The first two lines of the XCF GROUP display show values coded in the XCFGROUP. In this example, the default XCF group name is used, XCFMONITOR is set to 60 seconds, CONNECTTIMEOUT is 10 seconds, and RECOVERYTIMEOUT is 30 seconds. The LUNS LISTENER field is the IP address and port of the currently active LUNS. In the tabular portion, each row represents one member. It displays the member MVS system name, Telnet jobname, problem detection monitoring values, LUNS counter, LUNS configured type and rank, LUNS state, LUNS status flags, LUNR state, and LUNR status flags. The list is sorted by MVSNAME and then TNNAME.

At system RANS17, TLUNR1 and TLUNR2 both have connections with shared lu names allocated. TLUNR2 has a profile pending. The P flag will go off as soon as the LUNS replies to the profile build request. TLUNS1 is acting as both an active LUNS and a LUNR. It was originally configured as a primary with a rank of 255. There are two LUNRs still attempting rebuild. You know the LUNRs do not have allocated shared LU names because the L status flag is off. Therefore, the LUNS is not required to stay in RECOVER state. All members have the same LUNS counter value of 102. If there were an XCF communication problem and one or more of the members was still at 101, the new LUNS would not go to ACTIVE state.

At system RANS18, two LUNRs can not get a LUNS connection and can not recover. The C flag under PDMON tells you each LUNR has exceeded the CONNECTTIMEOUT time and has automatically quiesced any port that has shared LU name groups defined. No new connection requests will be directed to these LUNRs by a distributor. The R flag under PDMON tells you each LUNR may have dropped connections with shared LU names which would turn off the L status flag and allow the LUNS to go to ACTIVE state.

At RANS19, two active LUNRs have allocated LU names in use. An obeyfile has just finished processing for TLUNS1. The LUNS is in QUIESCE state to allow LUNS updates. Shared LU group definitions cause the LUNR to move to START and try to connect to the LUNS, setting the C flag. The P flag indicates a new profile with shared LU groups is pending.

VARY LUNS commands

```
>>--VARY--TCPIP,tnproc,LUNS,-+-START--------+----<<
                            +-QUIESCE------+
                            '-RESUME-------'

>>--VARY--TCPIP,tnproc,LUNS,-+-INACT,luname-+----<<
                            '-ACT,luname---'
```


Five new VARY commands are introduced to support the LU name coordination function.

The usage of the LUNS START command is described on the next slide “Takeover and Recovery”. The LUNS QUIESCE/RESUME pair is used to remove a standby LUNS from recovery contention or allow the LUNS definitions to be updated by an OBEYFILE. The format of the commands is shown here.

The VARY LUNS INACT and ACT commands inactivate and activate LU names in the LUNS LU table. The commands are similar to Classic Telnet VARY INACT and ACT commands. The only difference is you must specify LUNS instead of Telnet after the Telnet proc name to direct the command to the LUNS LU table instead of the Classic or LUNR LU table.

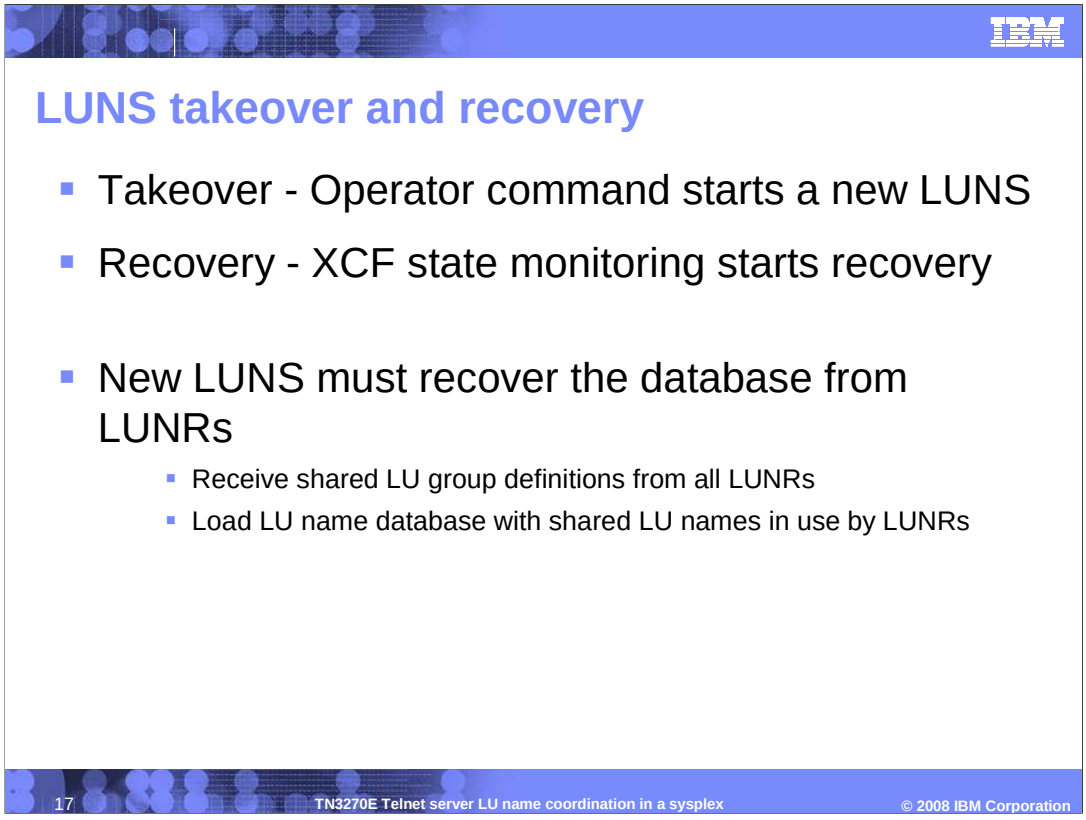

A new operator command is used to start a standby LUNS to become the active LUNS. Notification through XCF that a new LUNS is starting triggers the currently active LUNS to stop.

If an active LUNS fails, XCF notifies all group members of the failure and a backup LUNS automatically becomes active.

When a new LUNS starts, it must learn about the shared LU groups and already allocated LU names each LUNR owns before it can become active and start servicing the LUNRs.

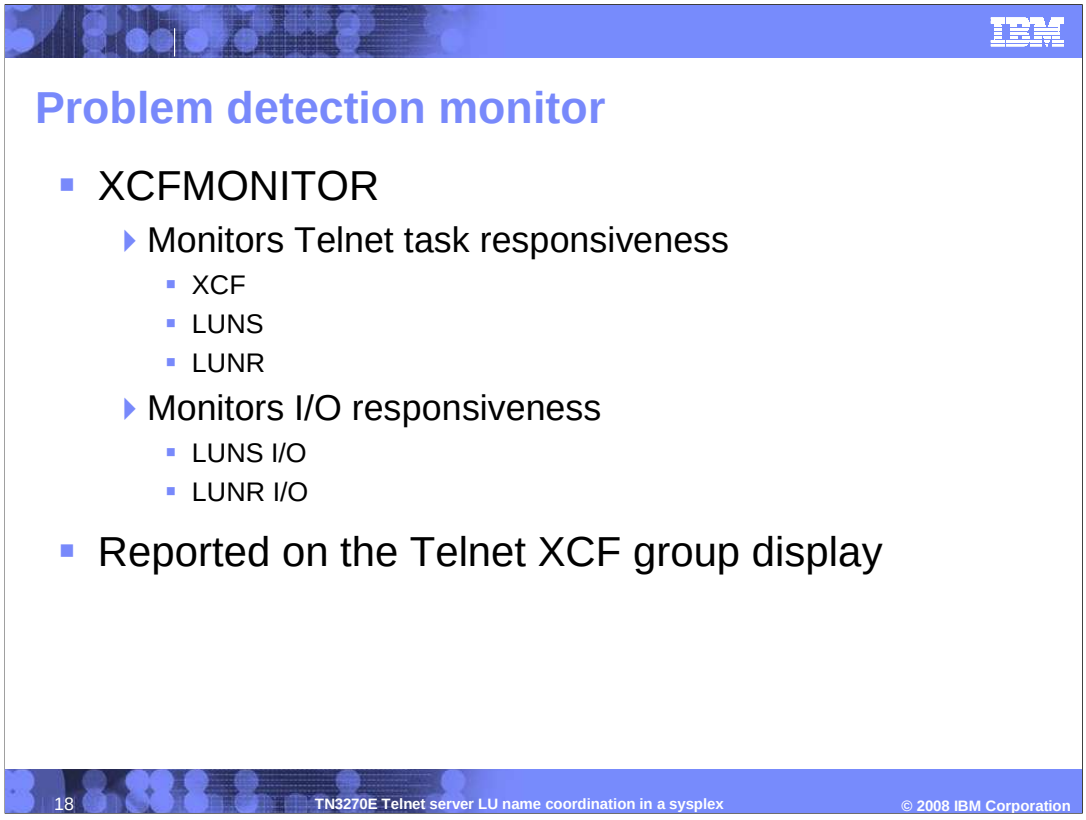

XCFMONITOR is a problem detection timer that monitors the health of the three Telnet tasks that represent the interface to XCF, the LUNS state manager, and the LUNR state manager. XCFMONITOR also monitors the send/receive processes on the administrative IP connections. When XCFMONITOR detects a problem, it is up to the system administrator to decide if action should be taken. At detection time, a message is issued indicating which of the five processes became unresponsive. If an I/O process problem is detected, the XCF STATS display might be helpful to determine what should be done next. If one of the three tasks has become unresponsive, the Telnet job probably needs to be recycled as soon as practical.

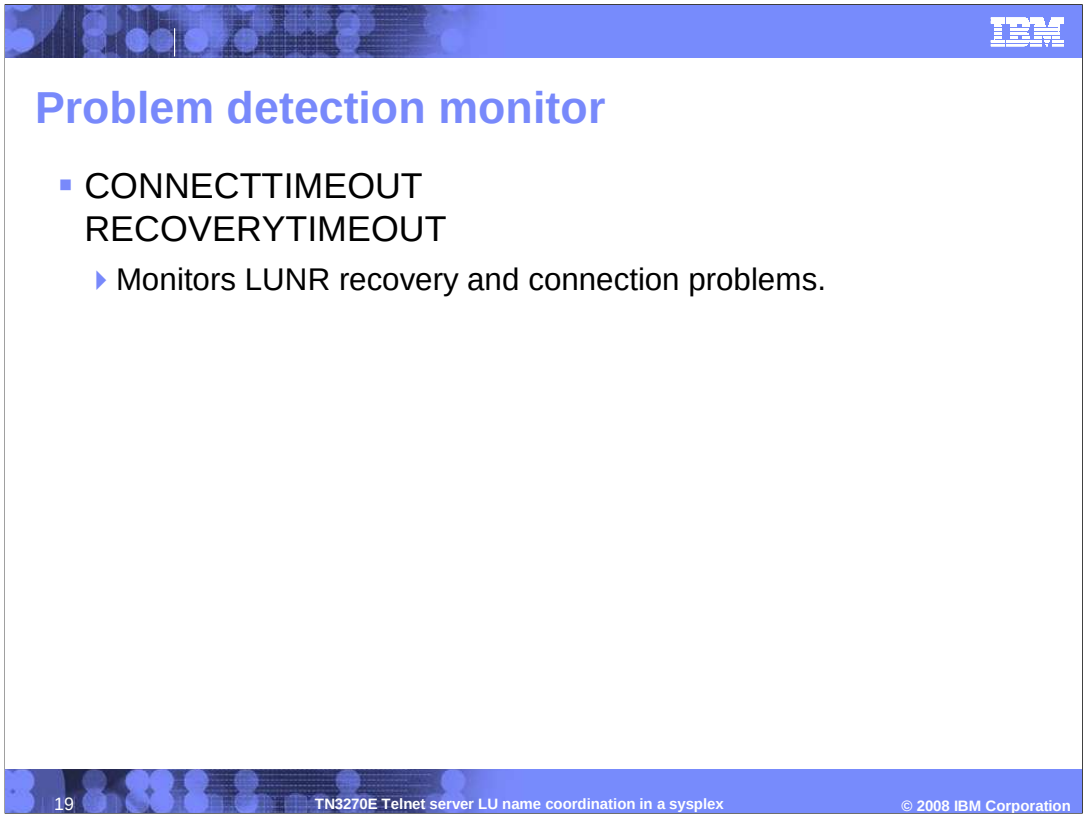

When RECOVERYTIMEOUT and CONNECTTIMEOUT detect a problem, action is taken automatically. RECOVERYTIMEOUT and CONNECTTIMEOUT are also part of the Problem Detection Monitoring system for XCF Telnet.

Problem detection monitor

- Reported on the Telnet XCF group display

```
PDMON
-----
XCR
```

  - X flag turned off if task or I/O becomes responsive again.
  - C,R flags turn off when a new LUNS connection is established.

Problem detection information can be displayed using the telnet XCF group display. There is a set of problem determination flags that can be displayed:

X indicates that the task or I/O has become unresponsive

C indicates that the LUNR has exceeded the CONNECTTIMEOUT time trying to establish an IP connection to the active LUNS. It has automatically quiesced any port that has shared LU name groups defined. No new connection requests will be directed to these LUNRs by a distributor.

R indicates that the LUNR has exceeded the RECOVERYTIMEOUT time trying to establish an IP connection to a recovering LUNS. It may have dropped connections with shared LU names.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM VTAM z/OS

A current list of other IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.