



IBM Software Group Enterprise Networking Solutions
z/OS® V1R11 Communications Server

FTP enhancements

z/OS Communications Server Development, Raleigh, North Carolina



© Copyright International Business Machines Corporation 2009. All rights reserved.

This presentation will give you an overview of the enhancements to the Communications Server in z/OS V1R11 for FTP.

Agenda

- FTP access to z/OS UNIX® named pipes
- FTP large-volume access
- FTP passive mode enhancements

In z/OS V1R11, FTP reads and writes files to and from z/OS UNIX named pipes.

FTP reads and writes extended format sequential data sets that have been allocated in the extended addressing space (EAS) of an extended address volume (EAV). This also includes support for data sets which are eligible for such an allocation (Format 8 DSCB).

The FTP client implements an option to allow the FTP client to simulate the use of extended passive mode, even in cases where the remote FTP server does not support EPSV.

Background: UNIX named pipes

- Visible in the file system with 'p' in file permissions

*file
size*



```

/SYSTEM/tmp ls -al
-rw-r----- 1 OMVSKERN OMVSGRP      39 Sep  3 22:22 eta
prwxr-x---  1 OMVSKERN OMVSGRP      0 Nov 18 14:07 eta.fifo
prwxr-x---  1 OMVSKERN OMVSGRP      0 Oct 30 14:46 eta1
-rw-r----- 1 OMVSKERN OMVSGRP      0 Aug  4 16:56 zero

```

Create it, list it, rename it, move it, delete it, chmod it!

A UNIX named pipe is a UNIX pipe identified by a pathname in the file system. It has characteristics of UNIX pipes and of UNIX files. Having a pathname, a named pipe is visible in the z/OS UNIX file system. The example on this slide is the output of the z/OS UNIX **ls** command. (The output has been edited to fit on this slide.) The files listed are a combination of z/OS UNIX named pipes and z/OS UNIX regular files.

When the first character of the permissions is 'p' as shown for the files eta.fifo and eta1, the file is the pathname of a named pipe. The other files, 'eta' and 'zero', are regular z/OS UNIX files.


The size of the named pipes is zero. This will always be the case, even when processes are actively writing to the named pipe.

You can create UNIX named pipes from the shell or from an application using the mkfifo command or the mkfifo() function call. The named pipe remains in place until you delete it. You can use the named pipe as many times as you like. A named pipe is not bound to a specific process, so any number of processes can use the named pipe. You can also create a named pipe using the mknod command or mknod() function call.

You can list, rename, move, delete, and chmod named pipes using UNIX shell commands just as you can regular files.

Background information: UNIX named pipe characteristics

- Opened for writing and reading at the same time
- Data is stored in a buffer
- FIFO – First In, First Out



The process that will write to the pipe and the process that will read from the pipe must have the pipe open at the same time. You cannot open a pipe for writing, write to it, and close it, expecting another process to come along at its leisure to read the pipe sometime later. The first process to open a named pipe blocks the `open()` call until the other process opens the other end of the pipe.


As a UNIX pipe, any data that a process writes to a named pipe is not stored in the file system; the contents of a pipe reside in a finite buffer of volatile storage. A UNIX pipe fills up when the pipe writer writes more data than the pipe reader can read. Writing to a named pipe that is filled blocks the pipe writer until space frees up in the storage buffer.

A UNIX pipe is empty when the pipe reader reads and processes data faster than the pipe writer produces it. Reading from a named pipe that is empty blocks the pipe reader until the pipe writer writes into the named pipe.

UNIX pipes are read and written strictly in FIFO (first in, first out) order. You cannot position the file marker to read in any other order. In fact, a common synonym for a named pipe is FIFO special file, or FIFO. Reading data from a named pipe removes the data from the pipe. The pipe reader cannot go back and read the same data from the pipe again. All writes to a named pipe are really appends to whatever is currently in the named pipe. It is impossible to replace the contents of a named pipe by writing to it.

FTP access to UNIX named pipes (also known as FIFOs)

- Saves total processing time when pre/post processing is needed for files transferred from/into z/OS
- Support available in both the z/OS FTP client and server
- FTP can be either the reading or writing end of the pipe
- PTFs for prior releases
- SAP



The z/OS FTP client and server are enhanced in z/OS V1R11 to write and read data directly into or from a z/OS UNIX named pipe.

The intention with this enhancement is better performance when files, for example, are transferred into z/OS and then need to be processed by a z/OS application. If that application supports reading from a z/OS UNIX named pipe, the transfer from the remote site can overlap processing by that z/OS application. That is, you have an unbroken pipe from the remote site that goes through a TCP connection and continues into a z/OS UNIX named pipe. When data is written from the remote client onto the TCP data connection, that data continues all the way into the post-processing application without further store and forward delays.

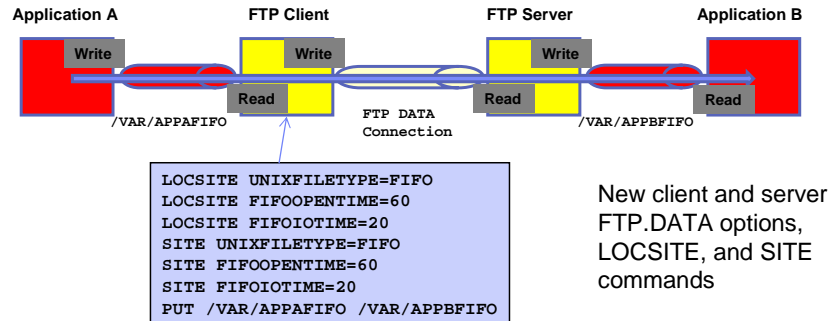
This z/OS FTP server support has been PTFed back to prior releases (z/OS V1R8, V1R9, and V1R10). The APAR number is PK71213.

Combined with the z/OS FTP Client API, some very interesting integrated solutions can be developed. One example is splitting large files into multiple streams and transmitting them over parallel FTP sessions for improved performance. Another example is collecting information from multiple sources on z/OS, building a stream of data, and sending it directly to a remote site without ever storing it on z/OS.

The SAP on DB2® for z/OS Unicode FASTLOAD conversion utility exploits named pipes. This product migrates SAP installations to UNICODE. You can learn more about this SAP product and how it exploits named pipes by consulting the document *SAP Unicode Conversion on DB2 for z/OS – Cookbook*. Joachim Rese, IBM Germany, 04 November 2008.

Example of z/OS FTP UNIX pipe support

- Example commands allow to 60 seconds to open the FIFO end points for read or write
- When all three “pipes” (/var/appafifo, the FTP data TCP connection, and /var/appbfifo) are open, transfer can begin
- When Application A writes a byte onto /var/appafifo
 - The byte arrives within milliseconds at Application B as data to be read over the /var/appbfifo
 - The pipe between Application A and Application B has with no store-and-forward in between



This example is not necessarily realistic, but it is chosen to explain all the elements of the FTP pipe support.

The UNIXFILETYPE option applies only to files in the z/OS UNIX file system. UNIXFILETYPE tells FTP if such a name is a UNIX file (UNIXFILETYPE FILE which is the default) or a UNIX pipe (UNIXFILETYPE FIFO).

When you code the UNIXFILETYPE FILE statement in FTP.DATA, you are configuring FTP for the prior release behavior. FTP can create, store into, or send data from z/OS UNIX regular files, just as it has always done. When UNIXFILETYPE is FILE, you cannot use FTP to store into or send data from named pipes.

Code UNIXFILETYPE FIFO to configure FTP to access named pipes. With the FIFO option, FTP can store into and send data from z/OS UNIX named pipes, but not regular files.

You can change this value for the current session with the site and locsite subcommands. If your client or server is dedicated to named pipes, code UNIXFILETYPE FIFO in FTP.DATA to set the value for all sessions. Otherwise, a practical guideline is to code UNIXFILETYPE FILE in FTP.DATA (or let it take the default value), and to set it with a SITE or LOCSITE subcommand to FIFO just before accessing named pipes.

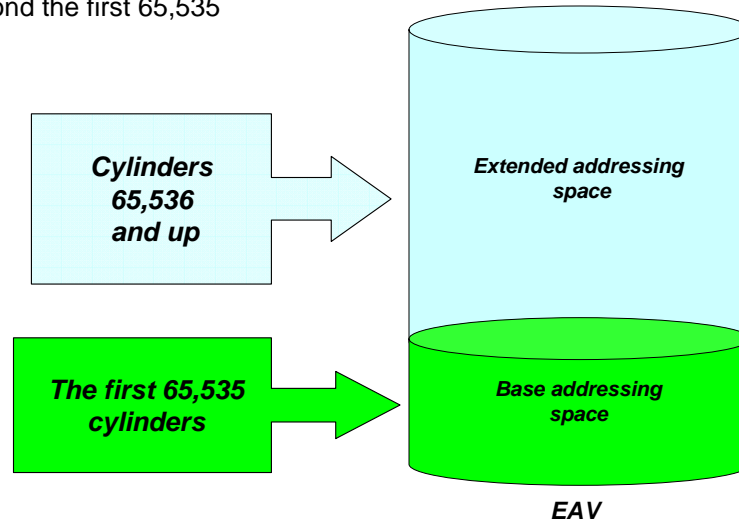
You cannot open a named pipe for writing until another process opens the named pipe for reading; you cannot open a named pipe for reading until another process opens the named pipe for writing. The first process to open the named pipe will block until the other process opens the named pipe.

You can control the amount of time z/OS FTP waits for the other process to open a named pipe by configuring the FIFOPENTIME value. When FTP opens a named pipe, it will wait up to FIFOPENTIME for another process to open the named pipe. If no other process opens the named pipe in FIFOPENTIME seconds, FTP will fail the file transfer.

When FTP reads from a named pipe, it will wait up to FIFOIOTIME seconds for the read to complete. If the read does not complete within FIFOIOTIME seconds (because the other process has not written any data for that amount of time), FTP will fail the transfer. When FTP stores into a named pipe, it will wait up to FIFOIOTIME seconds for the write to complete. If the write does not complete in FIFOIOTIME seconds, FTP will fail the transfer.

Background information – EAS of an EAV

- Extended Addressing Space (EAS) of an Extended Address Volume (EAV)
 - Cylinders beyond the first 65,535



The capacity of storage volumes has increased steadily over the years within the constraints of the 3390 storage device architecture. As of z/OS V1R10, z/OS has added support for DASD volumes having more than 65,520 cylinders. An Extended Address Volume (EAV) is a DASD volume that exceeds 65,520 cylinders. The DASD volume maximum for V1R10 is 262,668 cylinders.

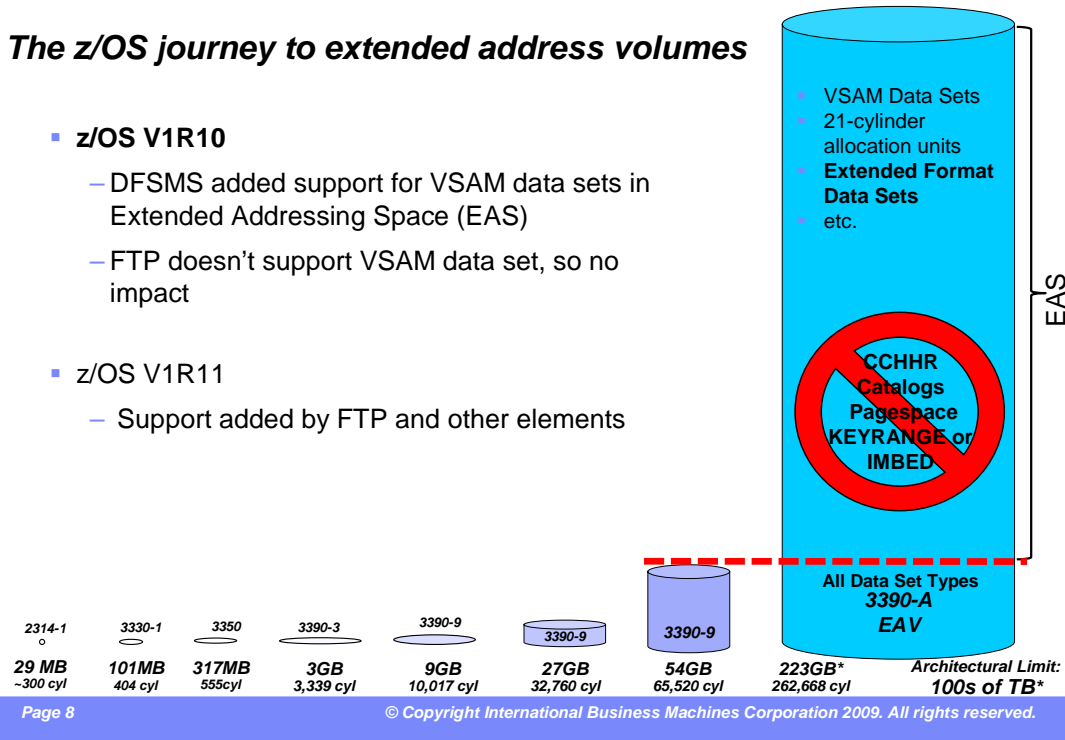
The EAS of an EAV is an abbreviation for Extended Addressing Space (EAS) of an EAV. It refers to those cylinders of a volume beyond the first 65,535.

In this diagram of an EAV, cylinders up to but not including cylinder 65,536 are in the base addressing space of the EAV. Cylinders starting with cylinder 65,536 are in the EAS of the EAV.

The z/OS journey to extended address volumes

- **z/OS V1R10**
 - DFSMS added support for VSAM data sets in Extended Addressing Space (EAS)
 - FTP doesn't support VSAM data set, so no impact

- **z/OS V1R11**
 - Support added by FTP and other elements



DFSMS is adding support over several releases for data sets in the extended addressing space (EAS) of extended address volumes (EAV). Initially, DFSMS only supported VSAM data sets in the EAS. In z/OS V1R11, DFSMS also supports extended format sequential data sets (not normal sequential data sets) in the EAS.

z/OS V1R11 adds some support for EAV extensions, including Extended Format sequential data sets and data sets spanning the 64K cylinder line. There is a new JCL keyword and data set attribute (EATTR). SMS, SMF, TSO/E, ISPF, and FTP support EAV extensions. Language Environment supports data sets with extended attributes, including those in the EAS on an EAV, for C/C++ programs. RACF® has discrete profile support for non-VSAM data sets in EAS. And SVCDUMP, SYSMDUMP, SNAP, ABDUMP, and IPCS support EAV extensions.

z/OS V1R11 support in FTP for extended address volumes

- DFSMS adds support for extended format sequential data sets eligible to reside in the EAS
- FTP adds support for reading/writing to/from existing EAS data sets, but not creating them (toleration mode)
 - FTP understands Format-8 DSCBs
 - FTP uses TRKADDR for track calculations
 - FTP qdisk option for SITE/LOCSITE output format changes

```
ftp> quote site qdisk
200-      Percent      Free      Free      Largest      Free
200- Volume Free      Cyls      Trks      Cyls-Trks  Exts  Use Attr
200- CPDLB3  45      1507      108      1440      2      22 Storage
200- CPDLB0  44      80486     156      461      0      25 Storage
200- CPDLB1  99      66619     5        65362    5      3 Storage
200 SITE command was accepted
ftp>
```

In z/OS V1R11, DFSMS adds support for extended format sequential data sets (not normal sequential data sets). Data sets that are eligible to reside in the EAS are allocated by DFSMS using a format-8 DSCB.

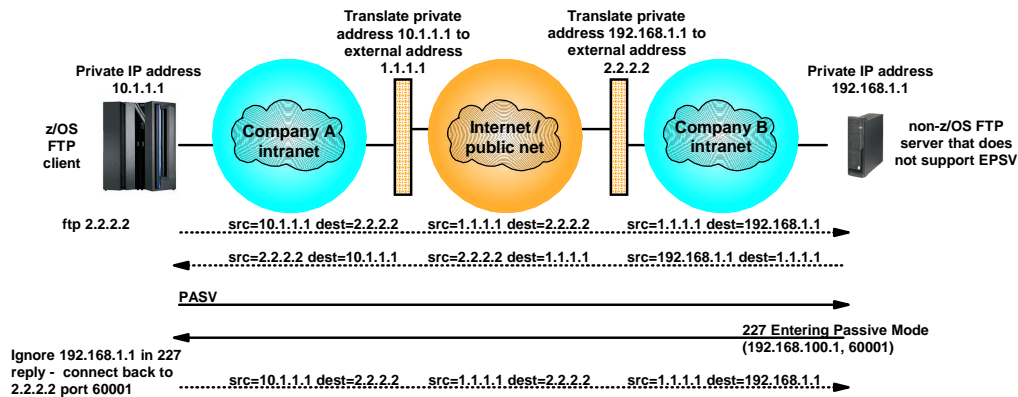
FTP reads and writes to such data sets in the EAS. FTP does not provide explicit configuration options to allocate data sets in the EAS or to specify data sets as eligible to reside in the EAS (unless implicitly done by way of SMS data class assignments).

FTP provides information about EAV volumes on the qdisk command. If you screen-scrape qdisk output, note that you will have to modify the screen-scraping logic since the qdisk output was changed between V1R10 and V1R11 to accommodate EAVs.

The LIST command output format did NOT change. However, if an output value is larger than the space allowed on the line, plus signs are displayed instead of the actual value.

FTP extended passive mode (even when servers do not know what it is)

- Extended passive mode FTP transfer solves a set of problems with FTP through NAT firewalls
 - But not all FTP servers support extended passive mode
- New z/OS FTP client support emulates extended passive mode behavior even when remote FTP server does not support EPSV



In z/OS V1R11, the z/OS FTP client implements an option that allows the FTP client to simulate the use of extended passive mode - even in cases where the remote FTP server does not support EPSV. The option enables passive mode data connections to succeed when the server is behind Network Address Translation (NAT) processing and might not advertise its correct external IP address to the client.

This option instructs the z/OS FTP client to ignore the IP address it receives in the PASV reply from the server. Also with this option, the z/OS FTP client uses the same IP address for the data connection that was used for the control connection.

The option is `PASSIVEIGNOREADDR` in the `FTP.DATA` configuration file. When `PASSIVEIGNOREADDR` is enabled, the client will ignore the IP address presented in the reply and use the same address it used to log into the FTP server. For the port, it will still use the port received in the PASV reply. `PASSIVEIGNOREADDR` can be either `TRUE` or `FALSE`. By default it is `FALSE` to prevent migration problems.

Also, FTP supports `PASSIVEIGNOREADDR` and `NOPASSIVEIGNOREADDR` options on the `LOCsite` command. The default is `NOPASSIVEIGNOREADDR`.

The `FWFRIENDLY` parameter must be set to `TRUE` for the `PASSIVEIGNOREADDR` option to have an effect.

These enhancements generally allow passive mode FTP data connections to be established through NAT firewalls in cases in which EPSV (extended passive mode support) established connections. In some cases, these enhancements might not be sufficient to allow the connections to be established through the NAT firewalls. These cases are when the FTP control connection is secured with SSL/TLS and the NAT firewalls (in addition to performing NAT) also implement dynamic filters based on the content of the PASV (or EPSV) reply. In these cases, use the FTP `CCC` command (clear command channel).



Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2 RACF z/OS

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

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.

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.

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.