



BROCADE FABRIC OS: LINK INITIALIZATION AND FILL WORD PRIMITIVE CONFIGURATION ON 8G FC PLATFORMS

Revision 1.0

July 26, 2010

Introduction

This document describes the Fabric OS (FOS) link initialization and fill word primitive configuration for Brocade 8G Fibre Channel platforms. By providing the flexibility to choose primitives for link initialization and fill word configuration, Brocade ensures full interoperability with all 8G capable FC HBAs and storage devices.

Note: The behaviors and configuration options described in this document only apply to 8G-capable ports.

Description

Up through the FOS v6.1.1 release, Brocade Fibre Channel switches used IDLE primitives both during link initialization and for fill words (denoted as "IDLE/IDLE"). This ensured successful link initialization between Brocade switch ports and end devices operating at 1G/2G/4G speeds.

With the introduction of 8G Fibre Channel, FC-PI-4 and FC-FS-3 standards stated that ARB should be used as fill words for lowering emission. To comply with the published FC standards, Brocade introduced options for ARB/ARB and IDLE/ARB link initialization/fill word support.

However, some 8G devices are not capable of properly establishing links with Brocade 8G Fibre Channel switches when ARB/ARB or IDLE/ARB primitives are used. These 8G devices require the legacy IDLE/IDLE sequence to achieve successful link initialization. To address this issue, Brocade has provided the ability to configure any of the three possible combinations (IDLE/IDLE, ARB/ARB, or IDLE/ARB) for link initialization and fill words. Any of these modes can be configured on an individual port basis using the CLI *portcfgfillword*. This CLI also provides an option to automatically switch from ARB/ARB to IDLE/ARB if the former is not able to successfully establish a link. (Please see the CLI documentation in the Appendix or in the FOS Command Reference Manual for further details). This capability is available in FOS v6.3.1 and later versions of firmware.

In summary, FOS v6.3.1 and later versions support the following modes for 8G F_Ports as a per port configuration.

MODE	MEANING
Mode 0	Use IDLE in link init and IDLE as fill word
Mode 1	Use ARB in link init and ARB as fill word
Mode 2	Use IDLE in link init and ARB as fill word
Mode3	Try mode 1 first; if it fails then try mode 2

Extended Fabrics/Long Distance E_Port Consideration

Standard E_Ports on Brocade's FOS-based platforms automatically set both the link initialization and fill word primitives to ARB regardless of the *portcfgfillword* setting when not in R_RDY mode.

In FOS v6.1.2 Brocade introduced support for configuring either ARB or IDLE primitives for fill words on Extended Fabrics E_ports when in VC_RDY flow control mode. The CLI *portcfglongdistance* was enhanced such that if the *vc_translation_link_init* parameter was configured as zero, then IDLE primitives were used as fill words. By default ARB primitives are used as fill words on long distance links.

When connecting to extension devices that do not support ARB primitives, *portcfglongdistance* *vc_translation_link_init* parameter should be set to zero.

It is important to note that some of the combinations of *vc_translation_link_init* value and fill word modes are incompatible. The following table provides the supportability matrix between *vc_translation_link_init* value and the fill word modes. **Users must not configure the unsupported values.**

Table 1:Extended Fabrics E_Port *vc_translation_link_init* and fill word Compatibility

<i>vc_translation_link_init</i> (Configured via <i>portcfglongdistance</i> CLI)	Fill word mode (Configured via <i>portcfgfillword</i> CLI)	Supported	LinkInit/Fill Word
0	0	Yes	IDLE/IDLE
0	1,2,3	No	NA
1	0	No	NA
1	1, 2, 3	Yes	ARB/ARB or IDLE/ARB

If there is a conflict between the fill word mode and *vc_translation_link_init* settings the *portcfglongdistance* command will warn the user of incompatibility, but it will not prevent the user from executing the CLI. That is, it will enable *vc_translation_link_init* as specified by the user, but it will NOT change the fill word mode.

For example, if the current fill word mode is 0 for the port 2/0, and if a user executes the command *portcfglongdistance* to set *vc_translation_link_init* to 1, then the command reports a warning message as follows:

```
> portcfglongdistance 2/0 LS,1 40
Warning: Port 16 portcfglongdistance vc_translation_link_init conflict with
portcfgfillword configuration.
```

Similar warning message is also displayed by the *portcfgfillword* CLI when it detects incompatibility with *portcfglongdistance* command settings. But this warning message from *portcfgfillword* is displayed only on FOS v6.2.1 or later releases on the FOS v6.2.x code stream and FOS v6.3.0 or later releases on the FOS v6.3.x code stream.

It should be noted that even though the warning message is displayed, the existing conflict between *vc_translation_link_init* and fill word mode setting will not be automatically resolved. The *portcfglongdistance* command executed above sets the *vc_translation_link_init* to 1, while the fill word mode remains configured to mode 0. Since this is an unsupported configuration, users must change the fill word mode to any of the compatible modes (1, 2 or 3) in order to ensure proper functionality. Users **must NOT** leave the port configuration in any of the unsupported modes listed in Table 1.

Verifying Current Link Init and Fill Word Mode Settings

Users can execute *portcfgshow* command to verify the configuration of fill word mode and VC Link Init settings for a given port.

```
sw0:root> portcfgshow 3/28
Area Number:                172
Speed Level:                AUTO (HW)
Fill Word:                0 (Idle-Idle)
AL_PA Offset 13:           OFF
Trunk Port                  ON
Long Distance               OFF
VC Link Init              OFF
Locked L_Port               OFF
Locked G_Port               OFF
Disabled E_Port             OFF
Locked E_Port               OFF
ISL R_RDY Mode              OFF
RSCN Suppressed             OFF
Persistent Disable          OFF
LOS TOV enable              OFF
NPIV capability              ON
QOS E_Port                  AE
Port Auto Disable:          OFF
Rate Limit                  OFF
EX Port                     OFF
Mirror Port                  OFF
Credit Recovery              ON
F_Port Buffers              OFF
NPIV PP Limit:              126
CSCTL mode:                 OFF
```

Procedure to Change Long Distance Fill Word Mode Settings

The following illustrates how a user can change the fill word setting of a long distance port using *portcfglongdistance* CLI, while ensuring that it does not conflict with the fill word mode setting that currently exists for that port.

Example:

Assume that a long distance port is currently configured to use IDLE/IDLE primitives. If a user wants to change the VC Link Init of this port to 1 (to use ARB primitives) then he/she can execute *portcfglongdistance* command with *vc_translation_link_init* parameter set to 1, and execute *portcfgfillword* CLI to set the fill word mode to 1, 2 or 3.

There is no specific sequence to follow to execute *portcfgfillword* and *portcfglongdistance* CLIs. These CLIs can be executed in any order but the end configuration should result in one of the supported configurations listed in Table 1. Users must not leave a port in any of the unsupported configurations.

Please use *portcfgshow* CLI to find out the fill word and VC Link Init values configured for a given port.

Behaviors in Different FOS Releases

The support for the ability to configure link initialization and fill word primitives has evolved over several FOS releases. The following table summarizes these changes.

FOS Release	Supported Combinations
Up through FOS v6.1.1	<p>Supports only IDLE/IDLE combination for F_Ports.</p> <p>Supports ARB/ARB for both normal and long distance E_Ports by default when in VC_RDY mode. But IDLE/IDLE combination is used on a long distance E_Port when that port is configured to use R_RDY flow control.</p>
FOS v6.1.2	<p>Supports the ability to configure either IDLE or ARB primitives as fill words on Extended Fabrics E_Ports via <i>portcfglongdistance</i> command and <i>vc_translation_link_init</i> parameter when in VC_RDY flow control mode. This enables support for frame based trunking on Extended Fabrics E_Ports in VC_RDY flow control mode irrespective of whether IDLE or ARB primitives are used for fill words.</p> <p>A new command <i>portcfgfillword</i> is introduced with options to configure either IDLE/IDLE by specifying mode 0 or ARB/ARB by specifying mode 1 (Note: This command was first introduced in FOS v6.2.0c and later supported in FOS v6.1.2)</p> <p>Supported modes via <i>portcfgfillword</i> CLI: Mode 0: IDLE/IDLE Mode 1: ARB/ARB</p>
FOS v6.2.0 FOS v6.2.0a FOS v6.2.0b	<p>8G switches are defaulted to use ARB/ARB primitives for 8G F_Ports. There is no command to change the combination to anything other than ARB/ARB on 8G F_Ports.</p> <p><i>Portcfglongdistance</i> command can still be used to change the fill words on Extended Fabrics E_Ports to use either IDLE or ARB primitives.</p>
FOS v6.2.0c FOS v6.2.0d FOS v6.2.1 and all FOS v6.2.1x patch releases FOS v6.2.2 and all FOS v6.2.2x patch releases FOS v6.3.0 and all FOS v6.3.0x patch releases	<p>A new command <i>portcfgfillword</i> is introduced with options to configure either IDLE/IDLE by specifying mode 0 or ARB/ARB by specifying mode 1. Users are recommended to upgrade to FOS v6.2.0d or later due to noise issues associated with mode 0 in FOS v6.2.0c.</p> <p>Supported modes via <i>portcfgfillword</i> CLI: Mode 0: IDLE/IDLE Mode 1: ARB/ARB</p> <p><i>Portcfglongdistance</i> command can still be used to change the fill words on Extended Fabrics E_Ports to use either IDLE or ARB primitives when in VC_RDY flow control mode. This enables support for frame based trunking on Extended Fabrics E_Ports in VC_RDY flow control mode irrespective of whether IDLE or ARB primitives are used for fill words.</p>

	<p>You must check for the supported compatibility (listed in Table 1) before using either <i>portcfgfillword</i> or <i>portcfglongdistance</i> commands on Extended Fabrics E_Ports.</p>
<p>FOS v6.3.1 and all FOS v6.3.1x patch releases</p> <p>FOS v6.4.0 and all FOS v6.4.0x patch releases</p>	<p><i>Portcfgfillword</i> CLI is enhanced to support additional modes 2 and 3. This allows any one of the combinations for a given port: IDLE/IDLE (mode 0), ARB/ARB (mode 1), or IDLE/ARB (mode 2)</p> <p>Mode 3 simply attempts mode 1 and then tries mode 2. The end result is either ARB/ARB or IDLE/ARB based on the end device configuration.</p> <p>Supported Modes: Mode 0: IDLE/IDLE Mode 1: ARB/ARB Mode 2: IDLE/ARB Mode 3: ARB/ARB then IDLE/ARB</p> <p><i>Portcfglongdistance</i> command can still be used to change the fill words on Extended Fabrics E_Ports to use either IDLE or ARB primitives when in VC_RDY flow control mode. This enables support for frame based trunking on Extended Fabrics E_Ports in VC_RDY flow control mode irrespective of whether IDLE or ARB primitives are used for fill words.</p> <p>You must check for supported compatibility (listed in Table 1) when using either <i>portcfgfillword</i> or <i>portcfglongdistance</i> commands on Extended Fabrics E_Ports</p>

Best practice guideline for 8G FC platforms.

For Brocade 8G platforms running FOS v6.3.1 or higher, configuring the fill word to mode 3 typically enables the port to use the right combination of Fibre Channel primitives to fully interoperate with the attached end device. Once set to mode 3, user intervention is no longer needed to configure the right primitives to enable the link between the Brocade 8G switch port and the end device.

If a device is known to require mode 2 (IDLE/ARB) then configuring mode 2 on the 8G switch port will activate the link sooner. If mode 3 is used in this environment, then the port will first attempt mode 1 then eventually settle to mode 2 which will cause a delay in initializing the port.

If a device requires IDLE/IDLE combination for link initialization then one needs to set the needed port to mode 0 using *portcfgfillword* CLI.

If you are changing the fill word primitives on FC long distance ports using *portcfglongdistance* CLI, please make sure that there is no conflict between the *portcfglongdistance* setting and *portcfgfillword* mode setting. Please refer to the **Procedure to Change Long Distance Fill Word Mode Settings** section of this document for additional details.

Firmware Migration Considerations

Please note that port link initialization and fill word configurations will not be impacted across FOS firmware upgrade/downgrade. However, please refer to appropriate release note documentation before performing any firmware upgrade or downgrade.

Appendix:

portcfgfillword CLI

Please note that the *portcfgfillword* CLI is not supported on releases prior to FOS v6.1.2 on FOS v6.1.x code stream and prior to FOS v6.2.0c on FOS v6.2.x code stream.

The description provided here is applicable to FOS v6.4.0 release. Please refer to the appropriate version of the Fabric OS Command Reference Manual or Release Notes for descriptions applicable to that release.

```
admin> help portcfgfillword
```

```
Administrative Commands                portCfgFillWord(1m)
```

NAME

```
portCfgFillWord - Configures the fill word for a single 8G
FC port.
```

SYNOPSIS

```
portcfgfillword [slot/]port, mode
```

DESCRIPTION

Use this command to configure the fill word of an 8G FC port. This command is not applicable to non-8G FC ports. This command disables and re-enables the port and the port comes online with the new fill word setting. The configuration is stored in nonvolatile memory and is persistent across switch reboots or power cycle.

NOTES

This configuration cannot be set on VE_Ports or VEX_Ports.

Use the `portCfgShow` command to display user-configured fill word settings.

The execution of this command is subject to Virtual Fabric or Admin Domain restrictions that may be in place. For details on command availability, refer to the Fabric OS Command Reference, Appendix A.

OPERANDS

This command has the following operands:

`slot` For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

`port` Specifies the number of the port to be configured, relative to its slot for bladed systems. Use `switchShow` for a listing of valid ports.

`mode` Specifies the fill word for portnumber. This operand is required. Valid values are one of the following:

```
0 | -idle-idle
    Sets IDLE mode in the Link Init and IDLE as fill
    word (default).
```

- 1 | -arbff-arbff
Sets ARB(ff) in the Link Init and ARB(ff) as fill word.
- 2 | -idle-arbff
Sets IDLE mode in the Link Init and ARB(ff) as fill word (SW).
- 3 | /-aa-then-ia

Fabric OS 2010-02-23 1

Administrative Commands portCfgFillWord(1m)

Attempts hardware arbff-arbff (mode 1) first. If the attempt fails to go into active state, this command executes software idle-arb (mode 2). This is the preferred mode.

EXAMPLES

To set the fill word of a port to ARBFF-ARBFF using the numeric mode notation:

```
switch:admin> portcfgfillword 2/3, 1
```

To set the fill word of a port to ARBFF-ARBFF using the -arbff-arbff operand:

```
switch:admin> portcfgfillword 2/3, -arbff-arbff
```

To set the fill word of a port to IDLE-ARBFF using the -idle-arbff operand:

```
switch:admin> portcfgfillword 2/3, -idle-arbff
```

To set the fill word of a port to the preferred mode (If HW arb-arbff fails, try SW idle-arbff):

```
switch:admin> portcfgfillword 27 -aa-then-ia
switch:admin> portcfgfillword 28 -aa-then-ia
```

```
switch:admin> portcfgshow 27
Area Number:      27
Speed Level:      AUTO(HW)
Fill Word:         2(SW Idle-Arbff)
(output truncated)
```

```
switch:admin> portcfgshow 28
Area Number:      28
Speed Level:      AUTO(HW)
Fill Word:         3(A-A then SW I-A)
(output truncated)
```

```
switch:admin> portcfgshow
Ports of Slot 0 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed           AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Fill Word       0 0 0 0   0 0 0 0   0 0 0 0   2 3 0 0 0
(output truncated)
```


portCfgLongDistance CLI

The following description of the *portCfgLongDistance* CLI corresponds to FOS v6.4 release. Please refer to the appropriate version of the Fabric OS Command Reference Manual for descriptions of this command applicable to that release.

Administrative Commands portCfgLongDistance(1m)

NAME

portCfgLongDistance - Configures a port to support long distance links.

SYNOPSIS

```
portcfglongdistance [slotnumber/]portnumber  
    [distance_level] [VC_Translation_Link_Init]  
    [desired_distance]
```

DESCRIPTION

Use this command to allocate sufficient numbers of full size frame buffers on a particular port to support a specified long distance link. The port can only be used as an E_Port. Changes made by this command are persistent across switch reboots and power cycles. This configuration can be cleared but not set on VE/VEX_Ports.

Long distance configuration allows native FC ports to run WAN/LAN connections. It ensures that the full bandwidth of a link or trunk can be utilized for a particular long distance configuration. The receiving port must have sufficient buffers available, so that the transmitting port can stuff the link with enough frames to fill the entire length of the link. As the distance between switches and the link speed increases, additional buffer-to-buffer credit are required to maintain maximum performance. If a port is configured as a long distance port, the remaining ports of that port group could be disabled, fail to initialize, or move to "buffer limited" mode due to a lack of frame buffer credits.

NOTES

The number of credits reserved for a port depends on the switch model and on the extended fabric mode for which it is configured. Not all distance modes are supported by all platforms. For example, the FC10-6 only supports LO and LS up to 120 km at 10 Gbps. Supported distance can vary greatly, depending on switch platform and available buffers. Refer to the FOS Administrator's Guide for complete details on platform-specific buffer credit models, long distance mode support, and maximum distance supported for specific hardware configurations.

This command requires a Brocade Extended Fabric License.

The execution of this command is subject to Admin Domain and Virtual Fabric restrictions that may be in place. For details on command availability, refer to the Fabric OS Command Reference, Appendix A.

The long distance modes L0.5, L1, and L2 are no longer supported in v5.3.0 and later.

Fabric OS

2010-06-16

1

Administrative Commands

portCfgLongDistance(lm)

A long distance link can also be configured to be part of a trunk group. Refer to portCfgTrunkPort help for details.

When a port is configured as a long-distance port, the output of portShow and switchShow displays the long-distance level. Refer to portShow help and switchShow help for details.

The portCfgISLMode and portCfgLongDistance LE, LD, or LS levels only can be enabled at the same time. Such an ISL uses the R_RDY mode of flow control over the long-distance link. While using R_RDY mode flow control, an E_Port cannot form trunk groups of long-distance links even if the trunking is enabled. This feature is not backward compatible with firmware versions that do not support it.

Ctrl-D will cancel the configuration update.

OPERANDS

This command has the following operands:

slotnumber

Specify the slot number (for bladed systems only), followed by a slash (/).

portnumber

Specify the number of the port to be configured, relative to its slot for bladed systems. Use switchShow for a list of valid ports. This operand is required.

distance_level

Specify the distance_level as one of the following (the numerical value representing each dis-

tance_level is shown in parentheses):

L0 (0):

Specify L0 to configure the port as a regular port. A total of 20 full size frame buffers will be reserved for data traffic regardless of the port's operating speed. Therefore, the maximum supported link distance is up to 10 Km at 1 Gbps, up to 5 Km at 2 Gbps, up to 2 Km at 4 Gbps, and up to 1 km at 8 Gbps, respectively.

LE(3):

Specify LE mode to configure an E_port distance greater than 5km and up to 10km. A total of 5, 10, 20, or 40 full size frame buffers are reserved for data traffic at port speeds of 1 Gbps, 2 Gbps, 4 Gbps, or 8 Gbps, respectively. LE does not require an extended fabric license.

Fabric OS

2010-06-16

2

Administrative Commands

portCfgLongDistance(1m)

LD (5):

Specify LD for automatic long distance configuration. The buffer credits for the given E_port are automatically configured based on the actual link distance. Up to a total of 250 full size frame buffers are reserved for data traffic depending on the distance measured during E_port initialization. Fabric OS v6.1.1 or later supports up to 3000km at 1 Gbps, up to 1500 km at 2 Gbps, and up to 750 km at 4 Gbps and 8 Gbps. If a value for desired_distance is specified, it will be used as the upper limit to the measured distance.

LS (6):

Specify LS mode to configure a long distance link with fixed buffer allocation >10 Km. Up to a total of 1452 full size frame buffers are reserved for data traffic depending on the specified desired_distance value.

vc_translation_link_init

On switches running Fabric OS v6.2.0 or later, this parameter specifies the fill words used on long distance links. When set to 0, the link uses IDLE fill

words. When set to 1, the link uses ARB fill words (default). The IDLE fill word option is not compatible with QoS configured links and Credit Recovery enabled links. You must disable these features before configuring long distance IDLE fill words.

On switches running firmware versions earlier than Fabric OS v6.2.0, this parameter controls the long distance link initialization sequence. Specify 1 to activate the long distance link initialization sequence. Specify 0 to deactivate this mode. This operand is optional. When the command is run without specifying a value, 1 is assigned automatically for a long distance link in VC_RDY flow control. Otherwise, 0 is assigned. For a long-distance link not configured for ISL R_RDY mode, this parameter must be set to 1; otherwise, it must be reset to 0.

desired_distance

This parameter is a required when a port is configured as an LD or an LS mode link. In LD mode, the value of desired_distance is the upper limit of the link distance and is used to calculate buffer availability for other ports in the same port group. When the measured distance exceeds the value of desired_distance, this value is used to allocate the buffers. In this case, the port operates in

degraded mode instead of being disabled due to insufficient buffers. In LS mode, the actual link distance is not measured, instead the desired_distance is used to allocate the buffers required for the port.

EXAMPLES

To configure a switch port 63 to support a 100km link and be initialized using the long distance link initialization protocol:

```
switch:admin> portcfglongdistance 4/15 LS 1 100
switch:admin> portshow 4/15
portCFlags: 0x1
portFlags: 0x20001          PRESENT LED
portType: 1.1
```

```

portState: 2    Offline
portPhys: 4    No_Light
portScn: 0
portId: 013f00
portWwn: 20:3f:00:60:69:00:02:48
Distance: static (desired = 100 Km)
portSpeed: 2Gbps

0      Interrupts:          9          Link_failure: 0          Frjt:
0      Unknown:            0          Loss_of_sync: 0         Fbsy:

Lli:           9          Loss_of_sig: 9
Proc_rqrd:     0          Protocol_err: 0
Timed_out:     0          Invalid_word: 0
Rx_flushed:    0          Invalid_crc: 0
Tx_unavail:    0          Delim_err: 0
Free_buffer:   0          Address_err: 0
Overrun:       0          Lr_in: 0
Suspended:    0          Lr_out: 0
Parity_err:    0          Ols_in: 0
2_parity_err: 0          Ols_out: 0
CMI_bus_err:   0

```

SEE ALSO

configure, portCfgISLMode, portCfgTrunkPort, portCfgShow, portShow, switchShow

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 02/10

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, and Turbolron are registered trademarks, and Brocade Assurance, DCFM, Extraordinary Networks, and Brocade NET Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

