Build FC I/O Fabric Super Highways Using ISL Virtualization

David Lytle, BCAF
Principal Engineer / Global Solutions Architect
Brocade Communications
dlytle@brocade.com

March 13, 2014
Session Number 14769

Legal Disclaimer

All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.

Nothing in this presentation shall be deemed to create a warranty of any kind, either express or implied, statutory or otherwise, including but not limited to, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement of third-party rights with respect to any products and services referenced herein.

OF MAINFRAME DOMINANCE!

Happy Birthday System z!

- Mainframes are in use in 90% of the *Fortune* 1000 and a large percentage of midmarket companies"
  - IBM and the Computer & Communications Industry Association

…more than 70 percent of all corporate data resides on mainframes. Legacy systems represent ***trillions of dollars*** in assets…
  - z/Journal

- About 25 % of IBM's annual revenue comes from the sale of mainframes and associated products like storage systems, software and services!
  - NY Times

Now on to Super Highways

- On land a Super Highway is one that is designed for travel at high speeds, having more than one lane for each direction of traffic.
- When all goes well then a smooth and fast trip from the origination point to the destination can easily be made.

And all of us have had the experience of traveling a good road system like this one!

Super Highways

- But even the best engineered and created highway system can lead to snarls and slow traffic as congestion begins to overwhelm the ability of the system to allow going from one place to another easily.

If I were a betting man then I would bet that everyone in this room has seen this occur all too commonly!

Storage Network Super Highways

- Computer systems in some ways resemble highway systems as they are designed for data to travel at high speeds by using more than one lane for each direction of traffic.

- When all goes well then a smooth and fast trip from the origination point to the destination can easily be made by data.

And all of us try to build our storage networking fabrics so that we obtain the best possible performance!

Storage Network Super Highways

- But even the best engineered and created SAN can lead to snarls and slow I/O traffic as the SAN grows and changes and congestion creates backpressure that begins to overwhelm the ability of the SAN to allow going from one place to another easily.

Fabric 1

Fabric 2

Original SAN

Today's SAN

If I were a betting man then I would bet that many of you in this room have seen this phenomena in action!

Well, maybe not quite this bad!

Storage Network Super Highways

- To ensure that the data flow within a storage network node, or between nodes, is fast, smooth and uncongested as possible, as architects we need to **minimize path latency** and **control and balance I/O flow** across ISL links.

- Minimizing PATH LATENCY begins with making sure that Fan In – Fan Out ratios on a data path do not overwhelm the link (over-subscription).

- Also of concern is the switch node itself since a switching device can add from 2 to 100s of microseconds of latency delay.

Minimizing Path Latency

Within The Switching Components

- Brocade's unique "cut through" frame routing means that a full frame does not have to reside in switch memory before it gets passed along
- "Cut through" frame routing allows the Brocade switch to have a low average frame latency delay, averaging about 2 microseconds, for fibre channel data frames – especially significant with SSD devices!
- Reducing I/O path "Latency" provides improved I/O performance to devices like SSD!

The last of a frame's bits are leaving the HBA or the FICON adapter…

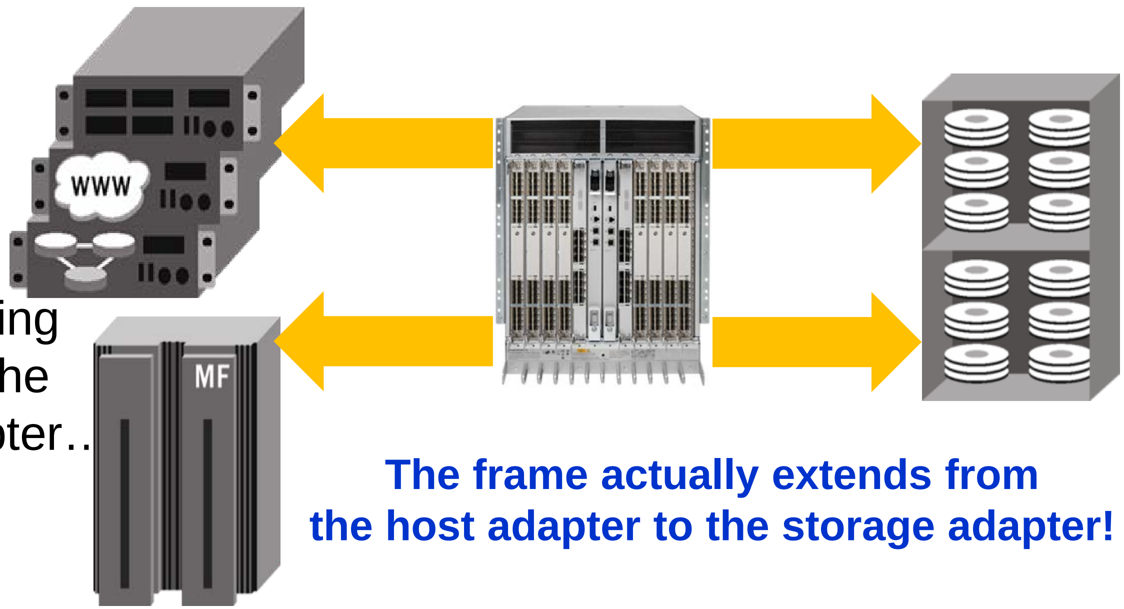

…as the first of a frame's bits are entering the storage port and the storage systems begins to work with that same frame.

The frame actually extends from the host adapter to the storage adapter!

Storage Network Super Highways

Between The Switching Components

- Controlling and balancing ISL flow is often not well understood but is often a major factor in poor performing storage networks.

- An inability to optimally use of all of the available ISL links is one major cause of poorly performing storage networks.

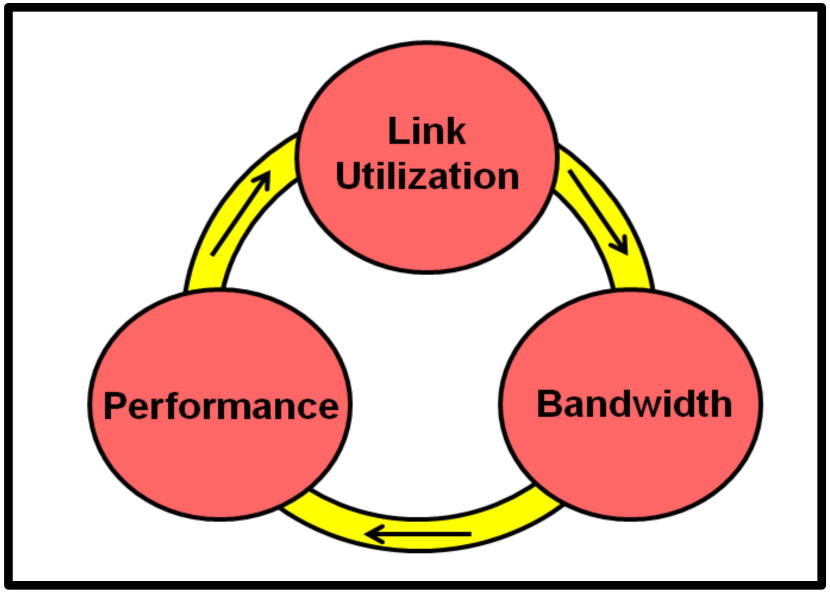

Storage Network Super Highways

Making Optimal Use of ALL ISL Links

- **Brocade supports 1 type of ISL hardware trunking and 3 types of ISL software Load Balancing**:

- **Brocade Trunkin**g allows blade ASICs to help optimize ISL links:
  - Hardware assisted trunking and requires specific configurations of ISL links
  - This is also hardware load balancing of the physical ISL links in a hardware trunk

- **Port-based Routing** allocates all ingress ports across one of the ISL links:
  - SID/DID software load balancing of physical ISL links and/or hardware trunks
  - Similar to the Fibre Channel protocol's standard shortest path first allocation

- **Device-based Routing** (at FOS 7.1) allocates ingress ports to ISLs:
  - Special SID/DID software load balancing of physical ISL links and/or hardware trunks
  - Will only assign ISL links to devices that need the use of an ISL

- **Exchanged-based Routing** (only for FCP for now)
  - SID/DID/OXID/RxPort software load balancing of physical ISL links and/or hardware trunks
  - Allocates the best ISL link to handle the data flow for a specific exchange

Device Based Routing (DBR) – a mainframe example

ISL Routes are only established when needed

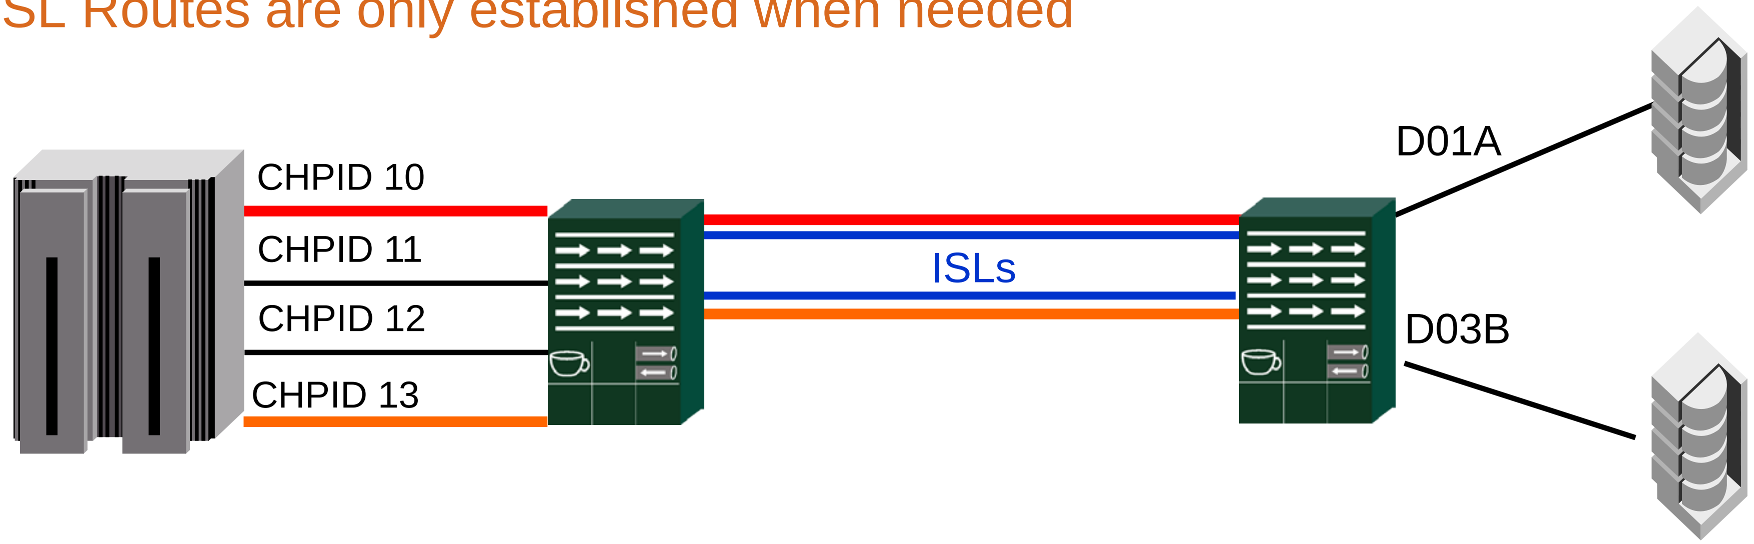

CHPID 10 logs in and then the first time CHPID 10 needs to send a frame to D01A an ISL route gets established.

CHPID 11 and 12 login, and are active, but they do not need to link to D01A or to D03B so they are not given a route across any ISL link. At this point only the 1st (top) ISL is in use.

The 2nd (bottom) ISL is the next route that gets assigned when CHPID 13 logs in and then does need to send a frame to D03B.

Since ISL routes are only established when needed, both ISLs are used for routes with data flows. Device ingress ports that will not utilize an ISL do not clutter up the allocation of ISLs.

Trunking with Device Based Routing (DBR) example

ISL Routes are optimized across the hardware trunks

- Port-based routing will assign I/O flows to each ISL of each trunk at PLOGI
- Device-based Routing (DBR) will modify I/O flows per exchange – as described
- Only F_Ports that require the use of an ISL are mapped to the ISL Links
- ISLs are never assigned to F_Ports that do not need the use of an ISL Link
- Load Balancing is based on: SID / DID
- DBR proactively ensures that each of the TRUNKS are evenly utilized for I/O flow
- The ASIC will proactively ensure that the ISL links of each trunk are optimized and avoid congestion

Port Groups

Each blade has port groups that are used to create the physical hardware trunking.

Then DBR makes sure that each trunk is optimized.

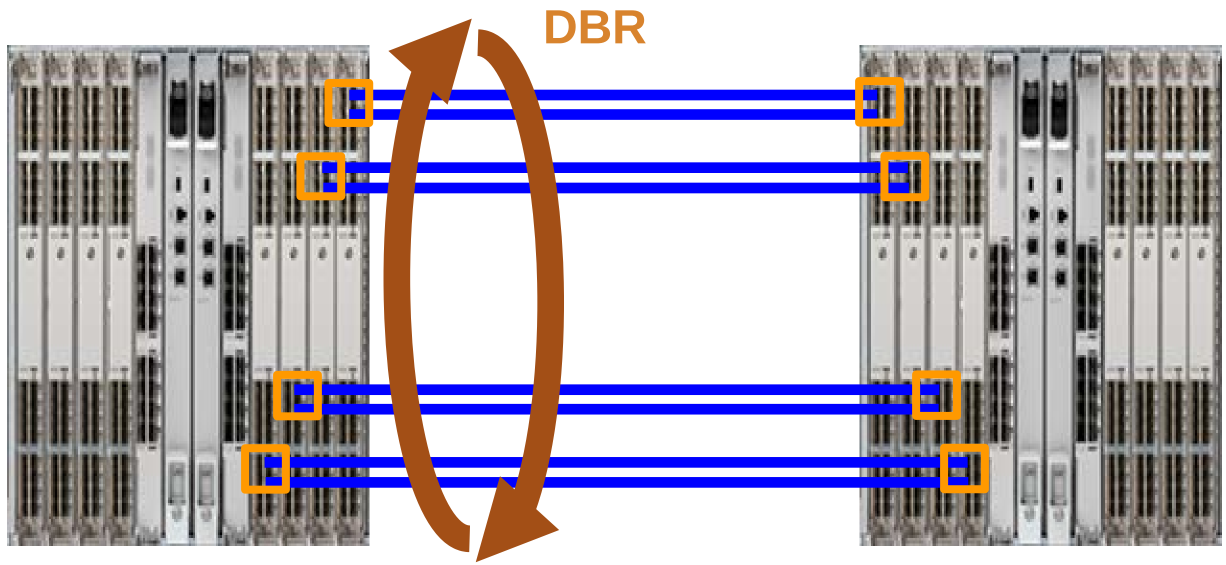

Storage Network Super Highways

- Even when ISL links are being used very well, **Head of Line Blocking** on an ISL link can create congestion and performance problems much like toll booths on a Super Highway can cause congestion and slow travel.

Head of Line Blocking
What Is the Problem?

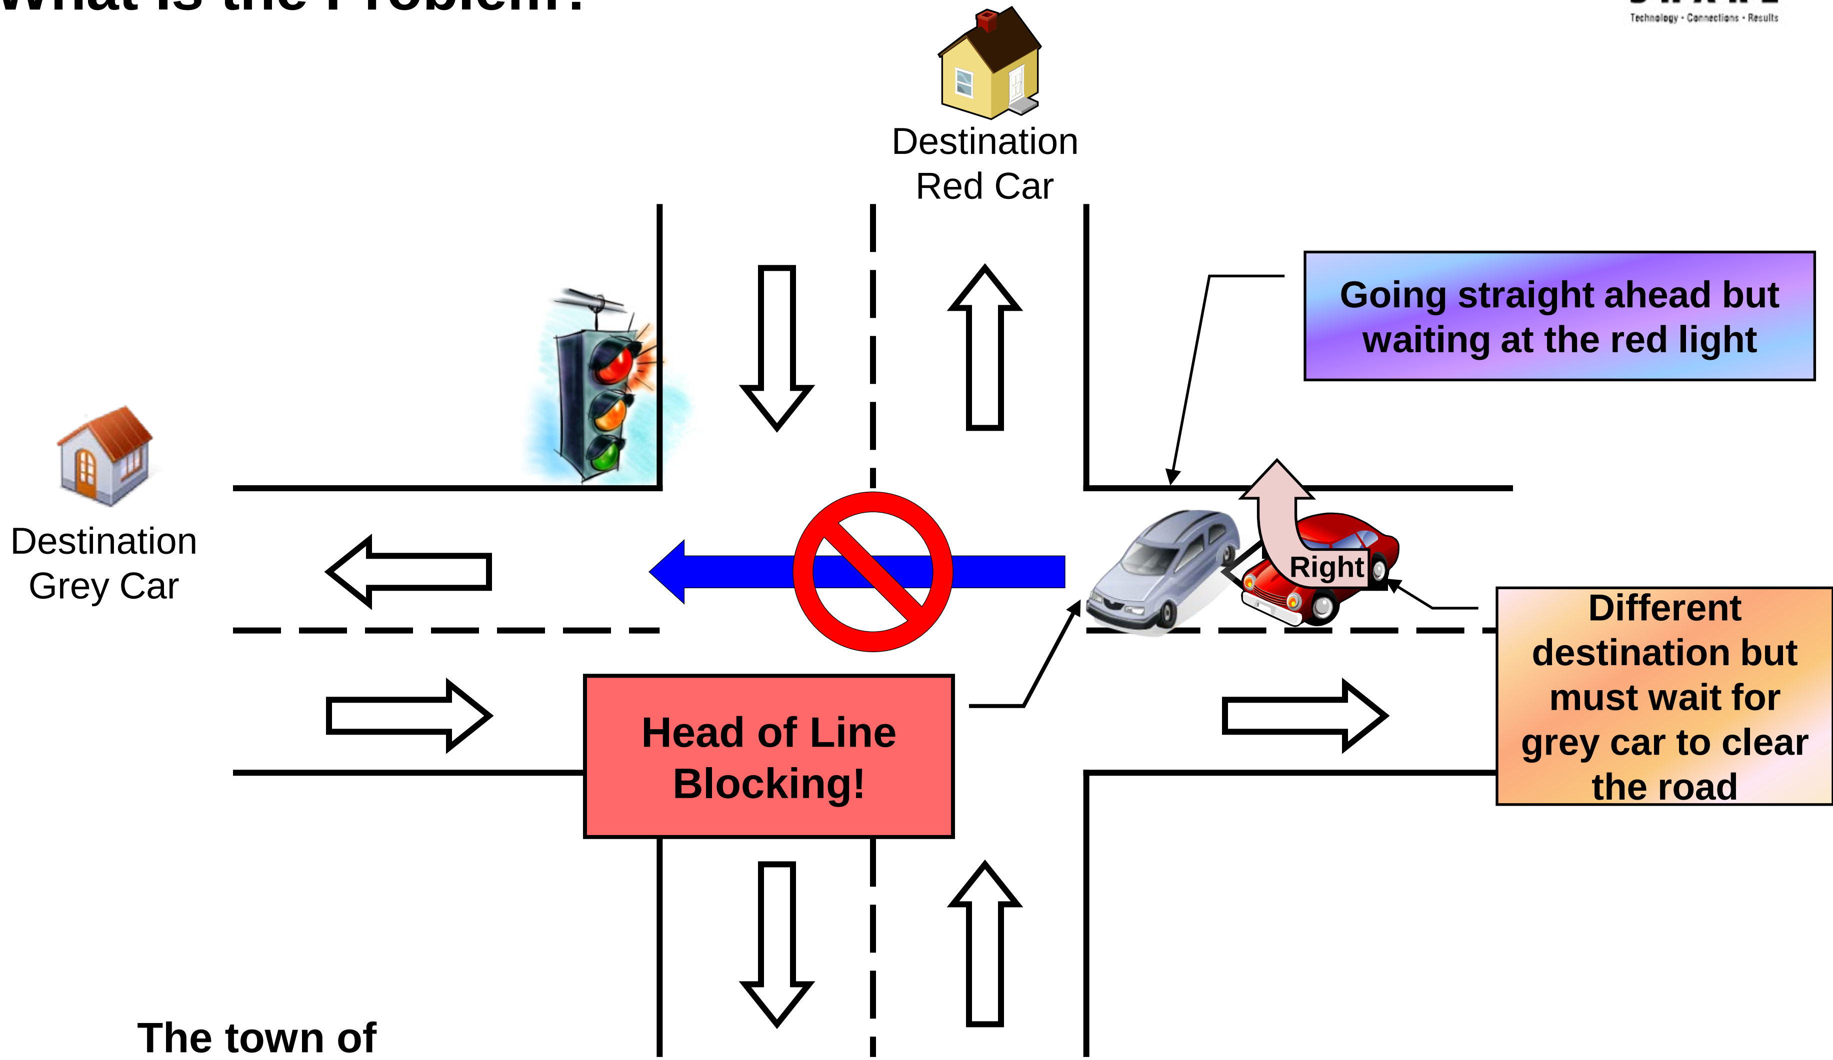

The town of "No Virtual Channels"

Roads are Full Duplex

Head of Line Blocking
What Is the Problem?

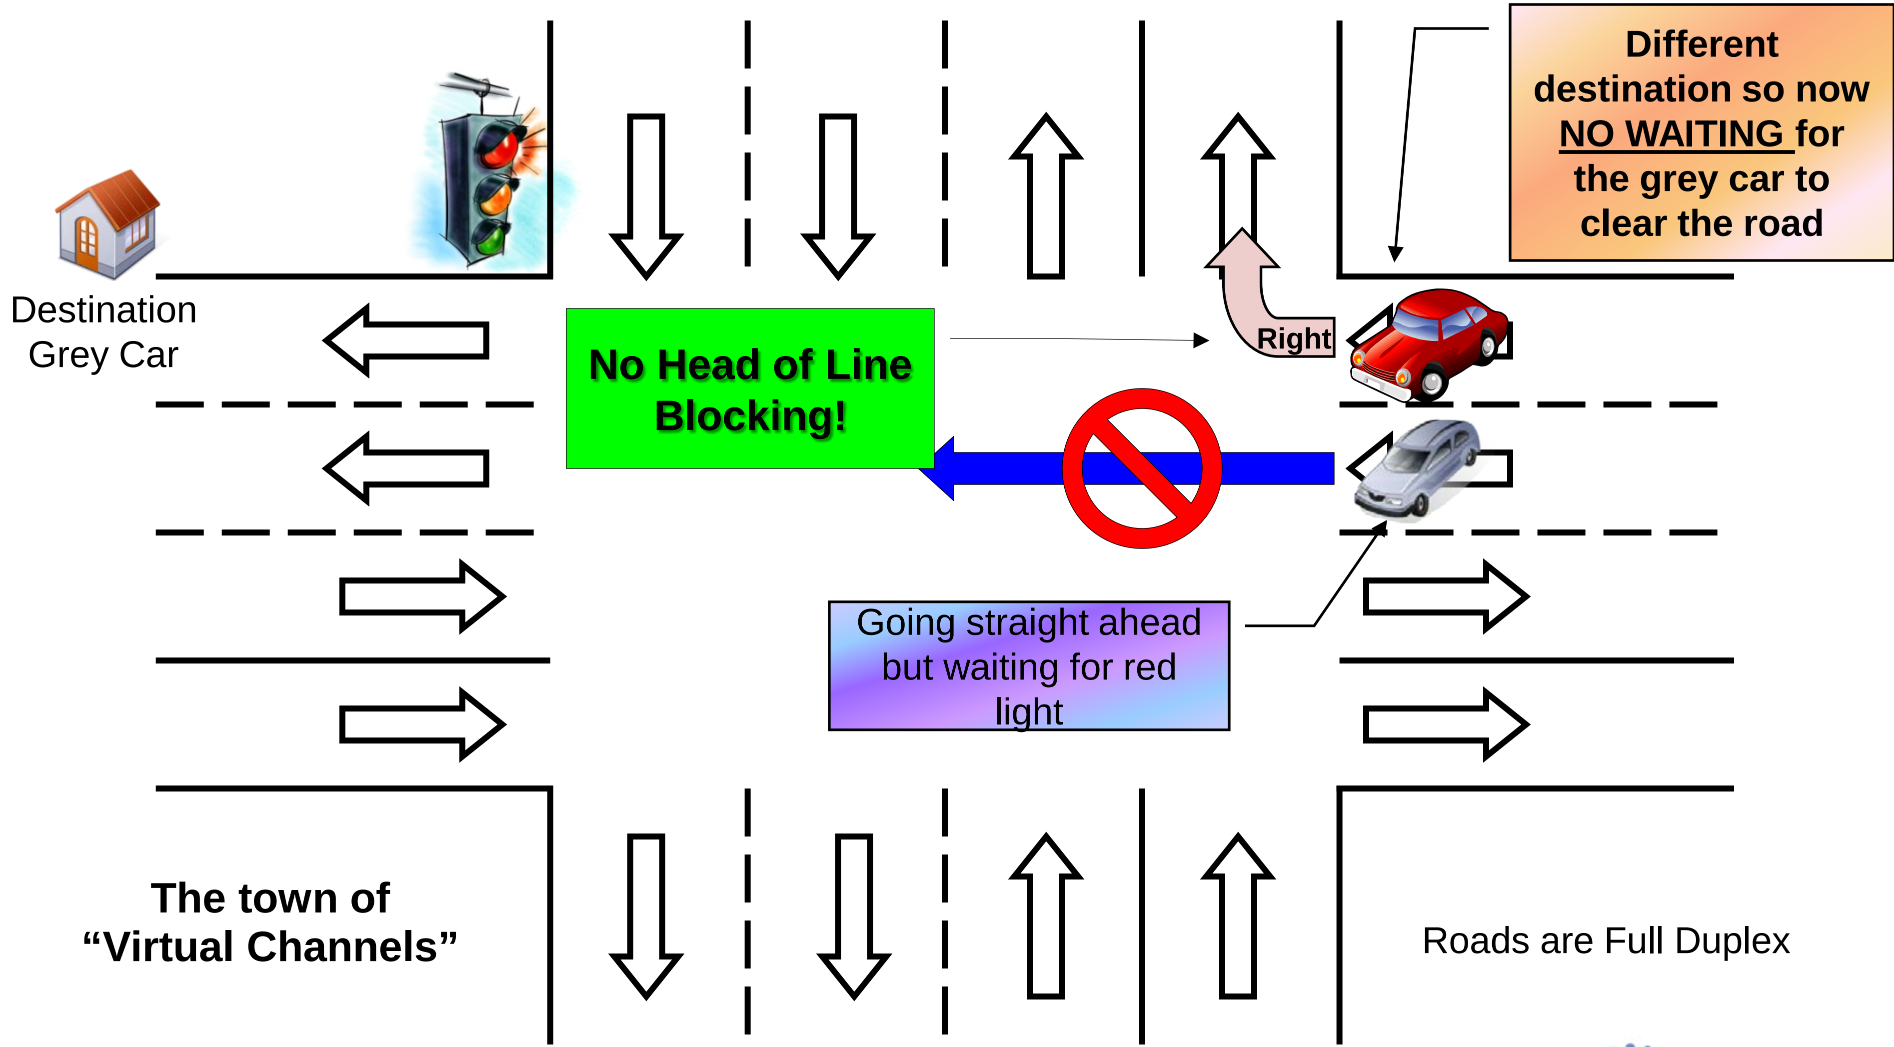

Head of Line Blocking

Needs to be avoided when possible

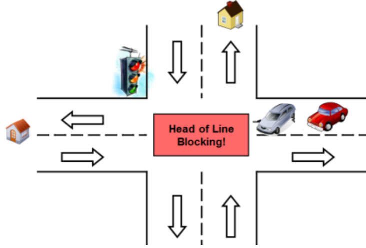

- As has been demonstrated, HoLB is a condition that users want to avoid whenever possible
- To minimize HoLB, a technical mechanism must be engineered and deployed into the architecture of the switching equipment that hosts your storage networks
- Brocade has discovered over its years of experience with Fibre Channel that it is best to adhere to the KISS principal
  - Keep It Simple Stupid
- Rather than building completely new technology, Brocade developed a way to segment the buffer credit (BC) pool on an ISL port and then direct the I/O frames to an appropriate BC segment on that ISL link
- Brocade calls this “Virtual Channels”

Virtual Channels (VCs)

Buffer Credit Pool segments established on each ISL link

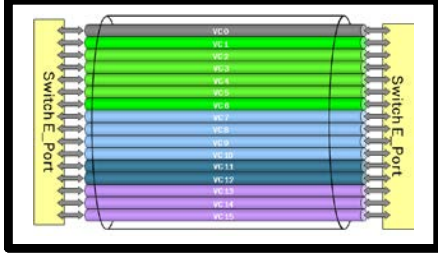

- Virtual Channels are a unique feature available on every Brocade 2, 4, 8 and 16Gbps fabric switch/Director.

- VC technology logically partitions the buffer credits on each ISL link into many different BC pools (virtual channels) and then prioritizes traffic to optimize performance and prevent head of line blocking.

- In addition, each VC has its own queues.

- *Of course an ISL is still just one fibre link so only a single lane (a single frame) of I/O traffic is passing across it, in each direction, physically at a time*

Virtual Channels help create a SAN Super Highway

Buffer Credit (BC) Pool segments provide multiple lanes of traffic on 1 link

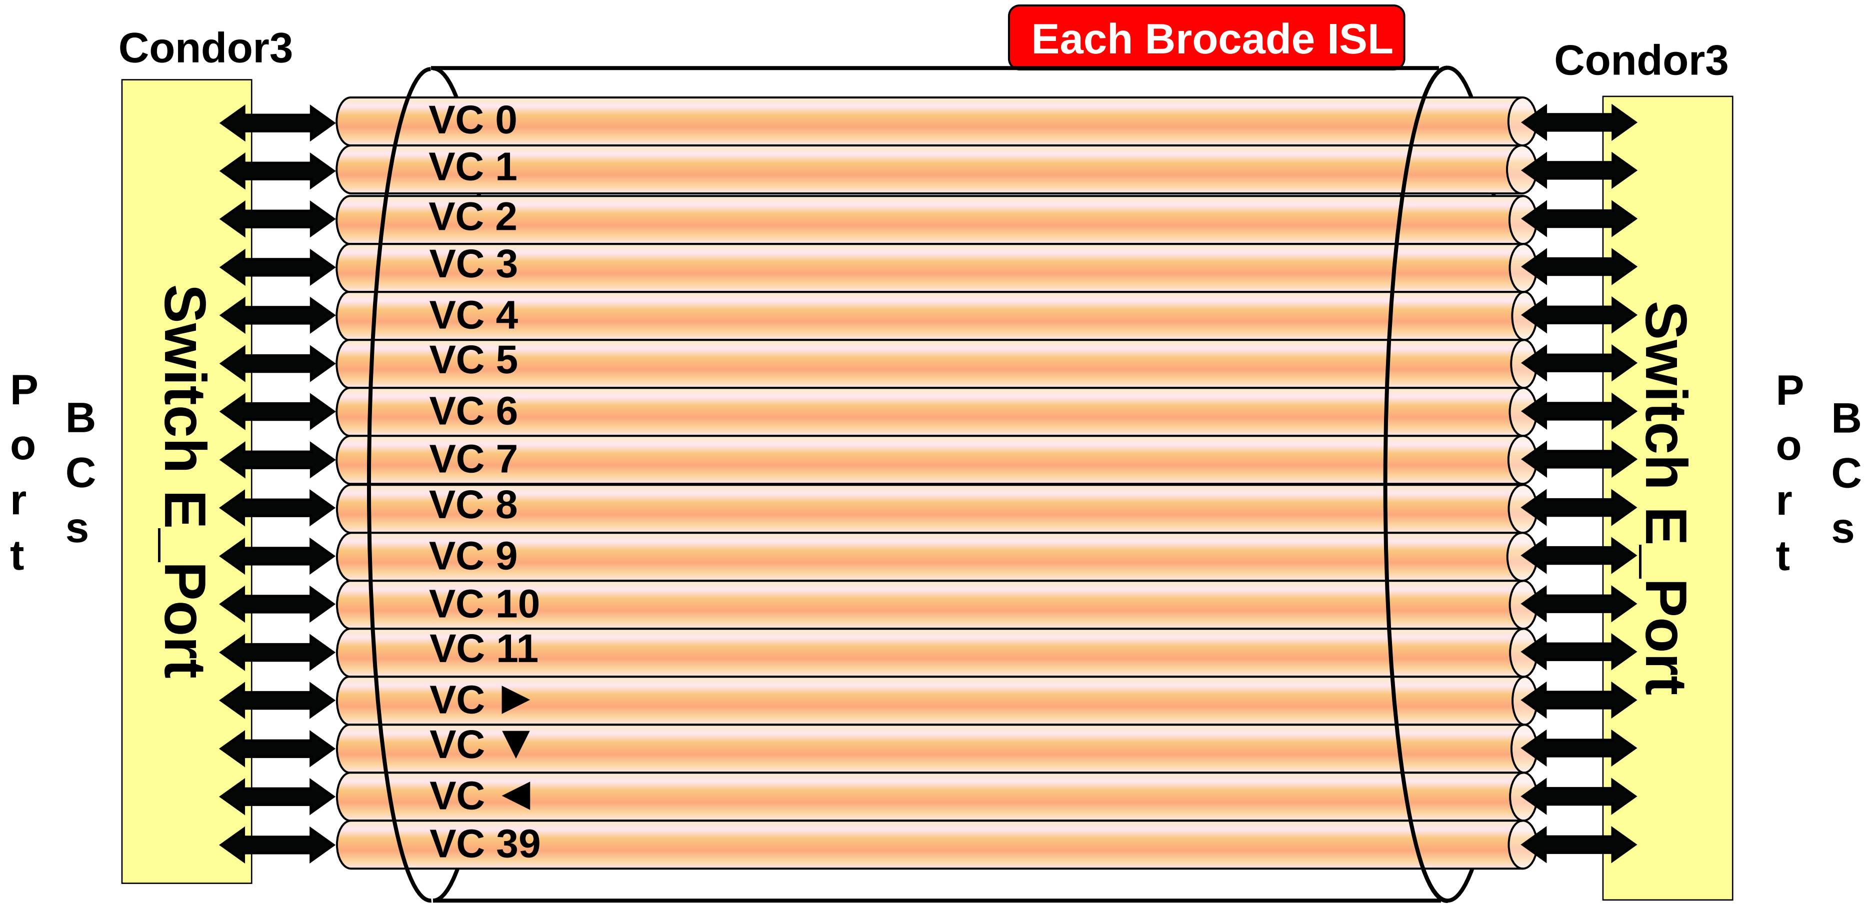

The Port BC pool is logically segmented into groups of virtual BCs (virtual channels) and then ingress ports can be assigned to a virtual channel segment if an ISL is required to be used.

Virtual Channels (VCs)

Some Concepts and Terminology

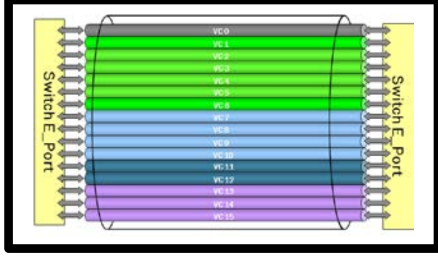

- VCs and Modes[1]:
  - ISL Links can be used for either local connectivity (e.g. core-edge topologies) or for long distance connectivity (more than a km or two).
  - Brocade uses “Modes” to differentiate the length of ISL links
  - Mode L0 is “Normal Mode” and is for local ISL link distances (<1 km)
  - Full Virtual Channel capability is <u>only available </u>in L0 mode! (Core-Edge)
- VCs and Quality of Service (QoS)[2]:
  - QoS is a tool for users to control end-to-end performance on a FC link
  - QoS allows users to choose high, medium or low performance for application I/O frames if/when congestion occurs in a SAN
  - QoS zones (the QoS tool) can be created for FCP but not for FICON
  - FICON is certified for QoS but it is not certified to use QoS zoning

Virtual Channels (VCs)

Some Concepts and Terminology

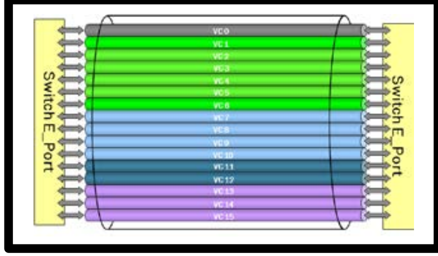

- VCs and Data Segments:
  - VC0 is used for class F I/O traffic service (F_Class service multiplexes frames at frame boundaries and is used for control and coordination of the internal behavior of the Fabric)
  - VC1 is used for Class 2 acknowledgements and Link Control
  - VC2-5 are used by F_Ports to send their data across an ISL link using “Medium” QoS performance parameters
  - VC6-7 are used for Broadcast and Multicast
  - VC8-14 are used for “Low” and “High” QoS performance for FCP I/O but not for FICON at all

- VCs and the F_Port FCID address:
  - The system uses the source Fibre Channel ID (FCID) address to direct its frames across an ISL link

Virtual Channel Allocations – without QoS

When QoS is not enabled

- The last nibble of the 2nd byte of the Fibre Channel ID (FCID) address is how an F_Port directs its frames across an ISL link
- In an F_Port 3-byte FCID of 61|0**3**|00, the **3** is the lower nibble of the second byte. This is used to find the appropriate ISL VC buffer credit segment to use to send an F_Ports frames across the ISL link.

Default ISL QoS not enabled				
VC	PID Used (DDAAPP)	Assigned to	Credits	Bandwidth
0	-	Class F	4	
1	-	Class 2 Ack/Link Control	-	
2	0, 4, 8, C	Data	5	100%
3	1, 5, 9, D	Data	5	
4	2, 6, A, E	Data	5	
5	3, 7, B, F	Data	5	
6	-	Class 3 Broadcast	1	
7	-	Broadcast/Multicast	1	

FICON or FCP with QoS disabled

Port Groups and VCs

Director Blades

Virtual Channel Allocations – with QoS

When QoS is enabled with or without QoS zones

- In an F_Port 3-byte FCID of 61|03|00, the **3** is the lower nibble of the second byte. This is used to find the appropriate ISL VC buffer credit segment to use to send an F_Ports frames across the ISL link.

QoS Enabled ISL				
VC	PID Used (DDAAPP)	Assigned to	Credits	Bandwidth
0	-	Class F	4	
1	-	Class 2 Ack/Link Control	-	
2	0, 4, 8, C	Data - Medium	2	30%
3	1, 5, 9, D	Data - Medium	2	
4	2, 6, A, E	Data - Medium	2	
5	3, 7, B, F	Data - Medium	2	
6	-	Class 3 Broadcast	1	
7	-	Broadcast/Multicast	1	
8	Round Robin*	Data - Low	2	10%
9	Round Robin*	Data - Low	2	
10	Round Robin*	Data - High	2	60%
11	Round Robin*	Data - High	2	
12	Round Robin*	Data - High	2	
13	Round Robin*	Data - High	2	
14	Round Robin*	Data - High	2	
15	-	-	-	

What FICON would do

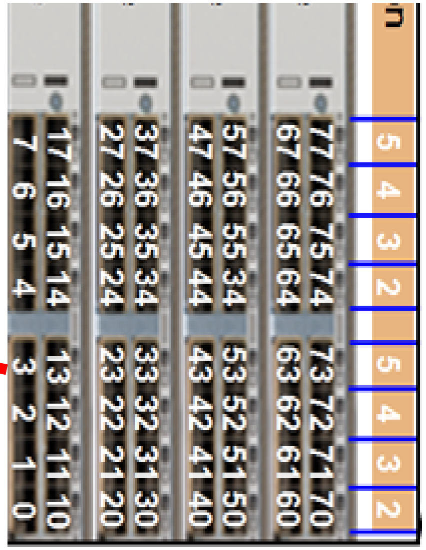

Port Groups and VCs

Director Blades

Only FCP can use QoS zones with QoS enabled

Virtual Channel Allocations – with/without QoS

There is a difference in the number of Buffer Credits

Default ISL QoS not enabled				
VC	PID Used (DDAAPP)	Assigned to	Credits	Bandwidth
0	-	Class F	4	
1	-	Class 2 Ack/Link Control	-	
2	0, 4, 8, C	Data	5	100%
3	1, 5, 9, D	Data	5	
4	2, 6, A, E	Data	5	
5	3, 7, B, F	Data	5	
6	-	Class 3 Broadcast	1	
7	-	Broadcast/Multicast	1	

- When QoS is DISABLED, more BCs are allocated to VCs 2-5
- 5 each per data VC

QoS Enabled ISL				
VC	PID Used (DDAAPP)	Assigned to	Credits	Bandwidth
0	-	Class F	4	
1	-	Class 2 Ack/Link Control	-	
2	0, 4, 8, C	Data - Medium	2	30%
3	1, 5, 9, D	Data - Medium	2	
4	2, 6, A, E	Data - Medium	2	
5	3, 7, B, F	Data - Medium	2	
6	-	Class 3 Broadcast	1	
7	-	Broadcast/Multicast	1	
8	Round Robin*	Data - Low	2	10%
9	Round Robin*	Data - Low	2	
10	Round Robin*	Data - High	2	60%
11	Round Robin*	Data - High	2	
12	Round Robin*	Data - High	2	
13	Round Robin*	Data - High	2	
14	Round Robin*	Data - High	2	
15	-	-	-	

- When QoS is ENABLED then the system must assume that QoS zones might get used and so it puts fewer BCs per VC
- 2 each per data VC

Virtual Channel Allocations – with QoS

FICON: Many buffer credits are not useable

QoS Enabled ISL			
VC	**PID Used (DDAAPP)**	**Assigned to**	**Credits**
0	-	Class F	4
1	-	Class 2 Ack/Link Control	-
2	0, 4, 8, C	Data - Medium	2
3	1, 5, 9, D	Data - Medium	2
4	2, 6, A, E	Data - Medium	2
5	3, 7, B, F	Data - Medium	2
6	-	Class 3 Broadcast	1
7	-	Broadcast/Multicast	1
8	Round Robin*	Data - Low	2
9	Round Robin*	Data - Low	2
10	Round Robin*	Data - High	2
11	Round Robin*	Data - High	2
12	Round Robin*	Data - High	2
13	Round Robin*	Data - High	2
14	Round Robin*	Data - High	2
15	-	-	-

FCP Can Use

Wasted For FICON!

- For FCP, QoS zones can be created that will make use of the High, Medium and Low data performance virtual channels:
  - VC2-5 – Medium I/O perf.
  - VC8-9 – Low I/O perf.
  - VC10-14 – High I/O perf.
- FICON only uses “Medium”

Virtual Channel Allocations in L0 mode

Assuming that all of the 32 ports are F_Ports

Virtual Channel Selection

PID	0	1	2	3	4	5	6	7
	8	9	A	B	C	D	E	F
VC	2	3	4	5	2	3	4	5

If we look at a 32 port blade this is what we will find:

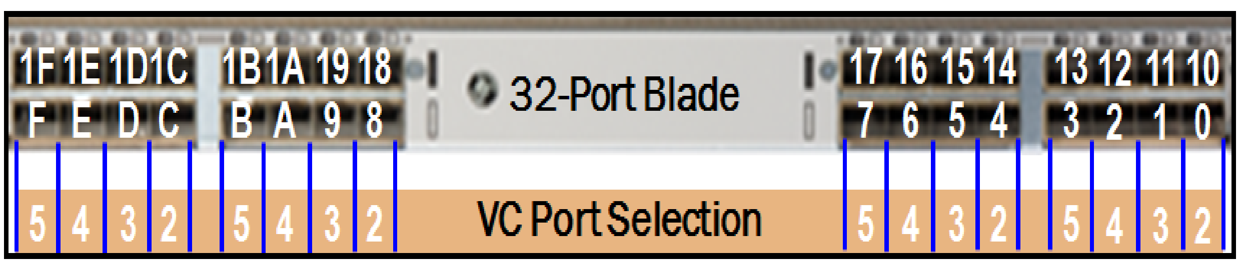

Port numbering is bottom-to-top or right-to-left

Virtual Channel Allocations in L0 mode

Assuming that all of the 48 ports are F_Ports

Virtual Channel Selection

PID	0	1	2	3	4	5	6	7
	8	9	A	B	C	D	E	F
VC	2	3	4	5	2	3	4	5

If we look at a 48 port blade this is what we will find:

48-Port Blade

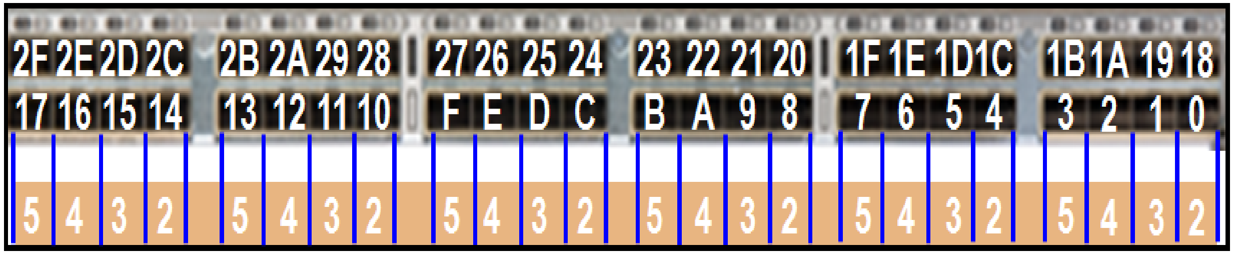

VC Port Selection

Port numbering is bottom-to-top or right-to-left

ISL Link Virtual Channels (VC)

Without QoS– L0 Long Distance Mode

- Disabling QoS prevents the use of VCs 8-14 on the ISL Links so that more BCs can be given to VC2-5
- Each and every ISL link (E_Port) will utilize VCs 0-7 only
- When QoS is not enabled, only eight Virtual Channels (out of the 15 useful VCs) are utilized on each of the user's ISL link(s)
- In L0 mode, without QoS, each of the VCs receives a predetermined number of buffer credits
- If more than 5 buffer credits per data VC (VCs 2-5) are required due to distance and/or average frame size, then make use of *portcfglongdistance* CLI command to deploy an adequate number for your environment using mode LE or LD or LS

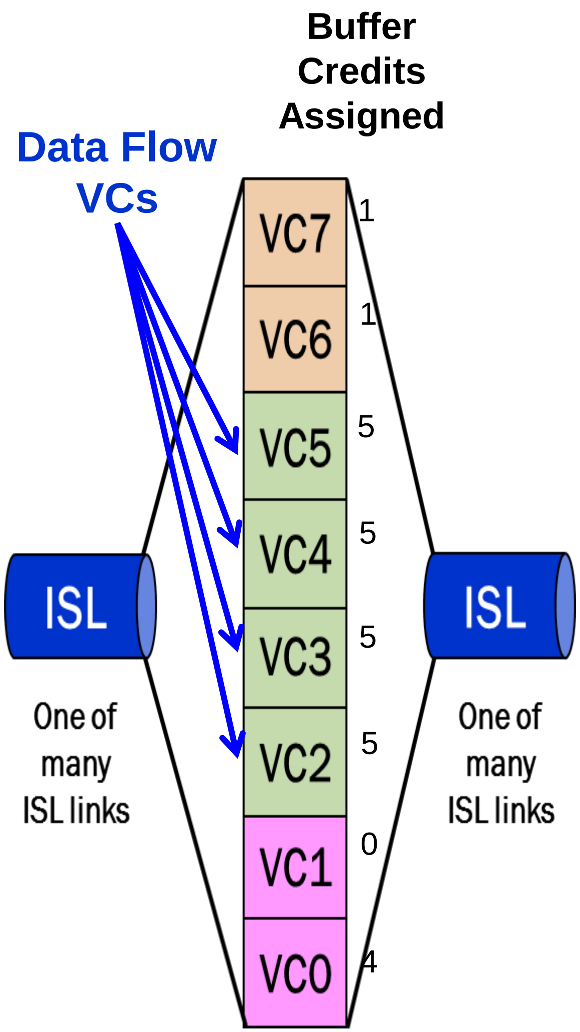

An ISL link can handle only one frame bi-directionally at a time but VCs minimize HoLB, maximizing ISL performance.

ISL Link Buffer Credits in Mode L0 – QoS Disabled

Physical BCs are mapped into Virtual Channel BCs

- Physical ISL ports have buffer credits allocated to them – the default is 26
- Depending upon link speed, this is for distances of 1 to 5 km

26 buffer credits on both of the ISL ports
- VCs 2-5 get 5 BCs each (20 total)
- VC 0 gets 4 BCs and VCs 6-7 get 1 BC each
- 0 shared BCs are available

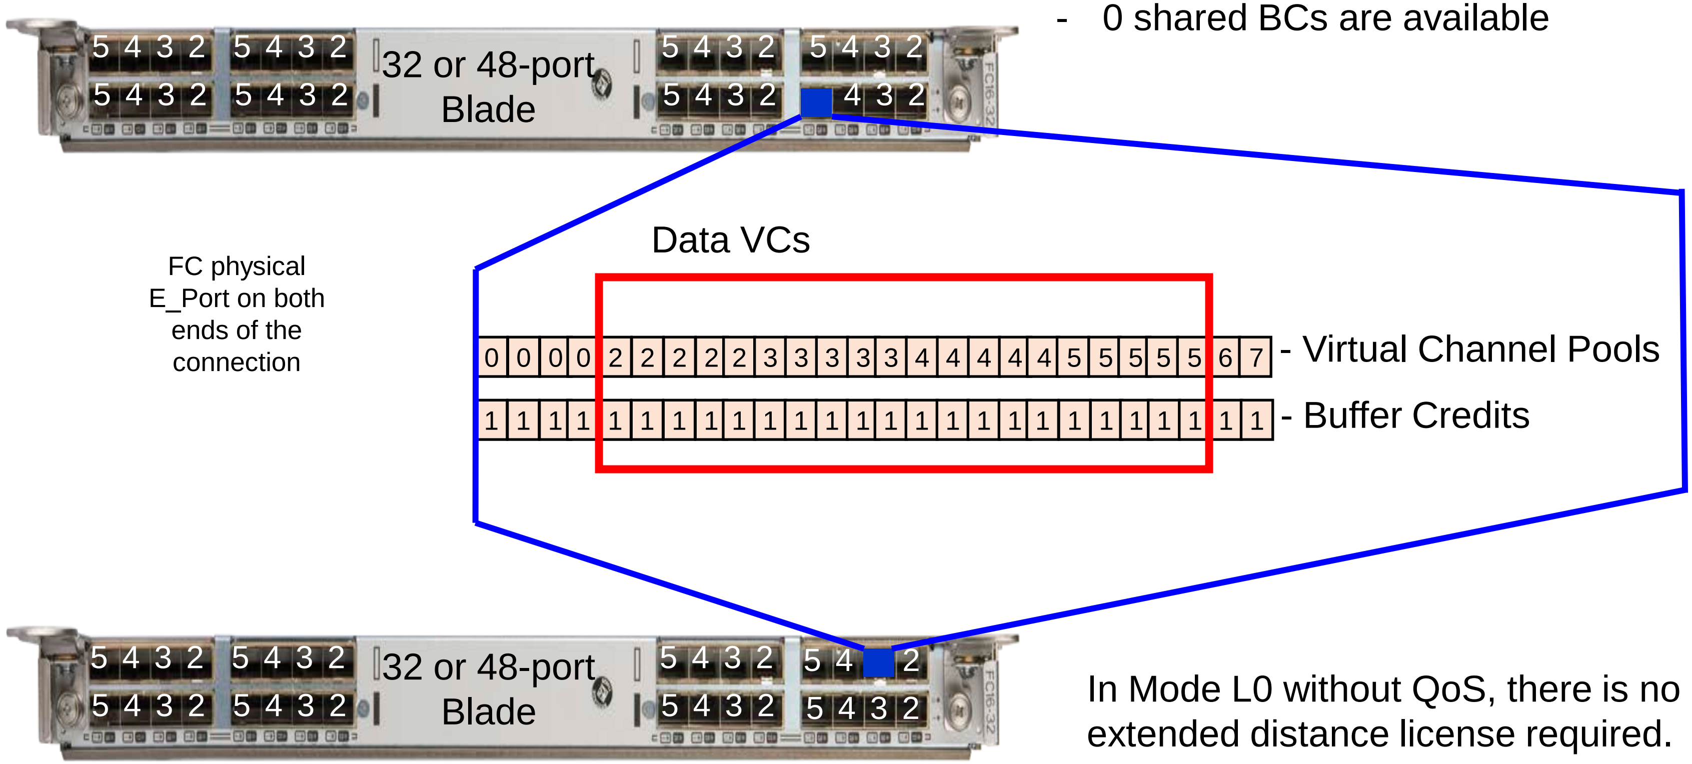

In Mode L0 without QoS, there is no extended distance license required.

ISL Link Buffer Credits in Mode L0 – QoS Disabled

Physical BCs are mapped into Virtual Channel BCs

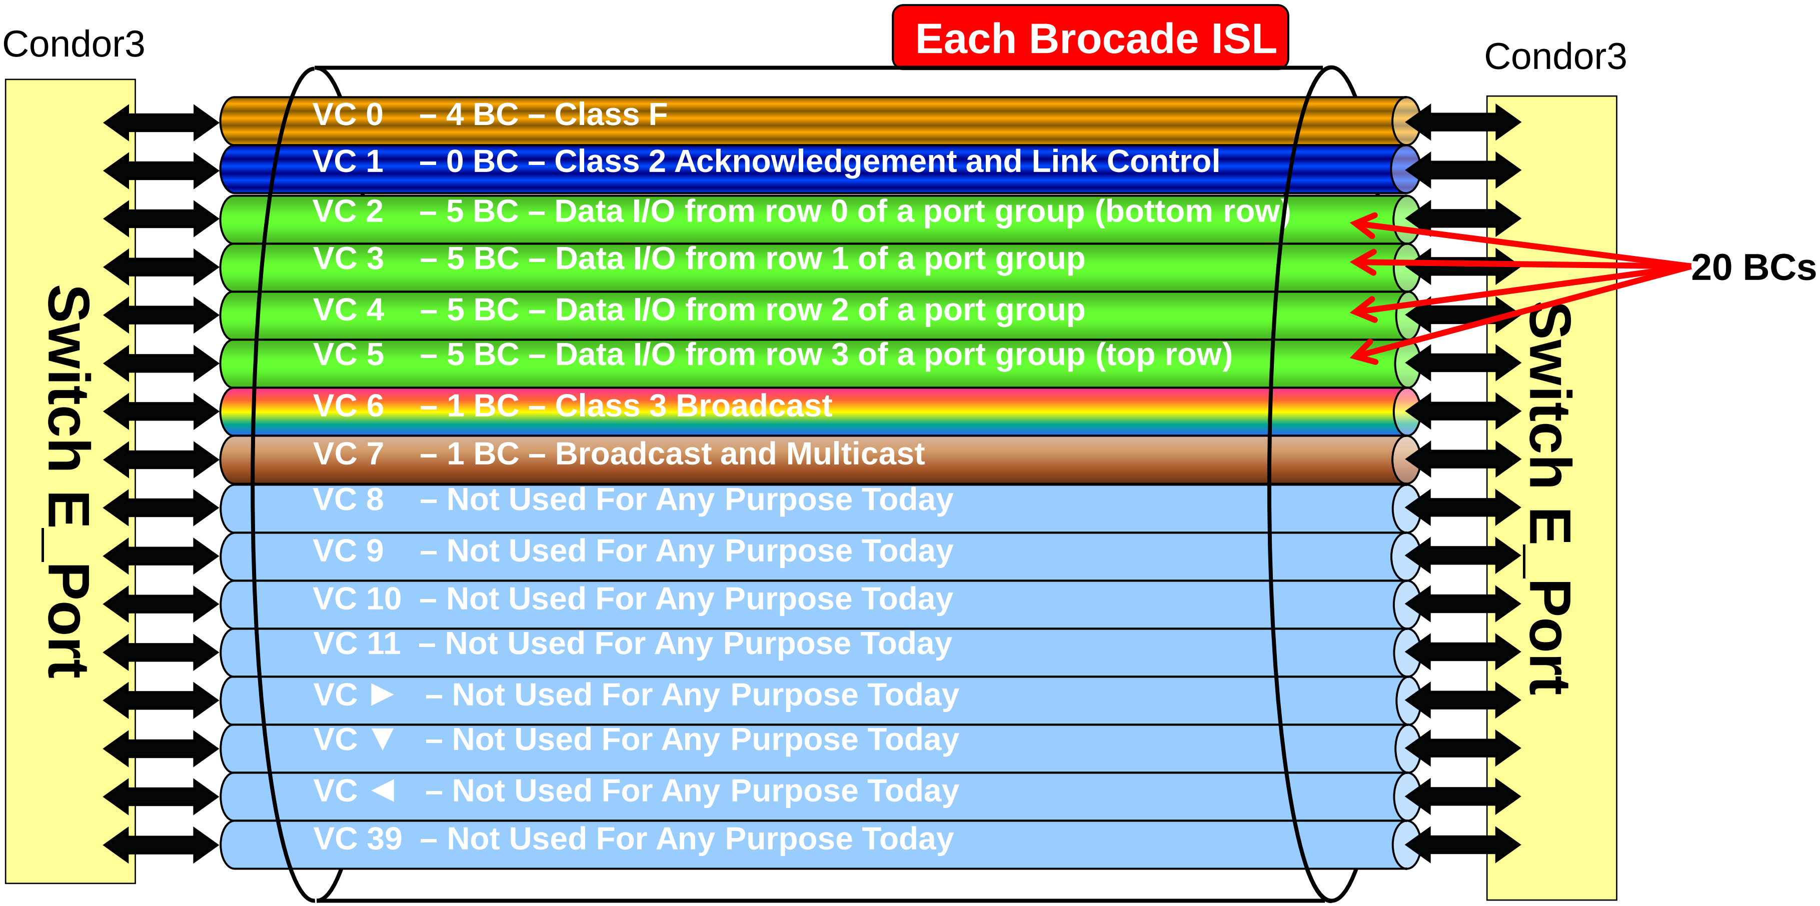

If QoS is disabled, FOS provides E_Ports with 20 BCs to be used for its data I/O across the VCs. The 20 BCs are divided evenly between VCs 2 through 5.

ISL Link Buffer Credits in Mode L0 – QoS Enabled

Physical BCs are mapped into Virtual Channel BCs

- With QoS enabled (the default for FICON), by default FOS provides 34 BCs to the virtual channel segments 0-14, on any given ISL link (L0 mode[1])
- Since FICON cannot use QoS zoning, the physical BCs that are provisioned out on VCs 8-14 are unavailable and therefore wasted for System z
- In L0 mode (<=1km distance) each of the VCs receives a predetermined number of buffer credits – but this number of BCs is a few less than when QoS is disabled because VCs 8-14 are given BCs that cannot be used for FICON
- Unless there are special circumstances, like a non-transparent DWDM as part of the long distance link, each and every ISL link (E_Port) should utilize Virtual Channels
- However, there are 6 BCs that are shared for VC 2-5 so those VCs could each "borrow" 1 BC if required from the shared pool
- If more than 2 or 3 BCs per data VC are required due to distance and/or average frame size, users at FOS 7.2 or later can use the *portcfglongdistance* CLI commands to deploy an adequate number for your environment

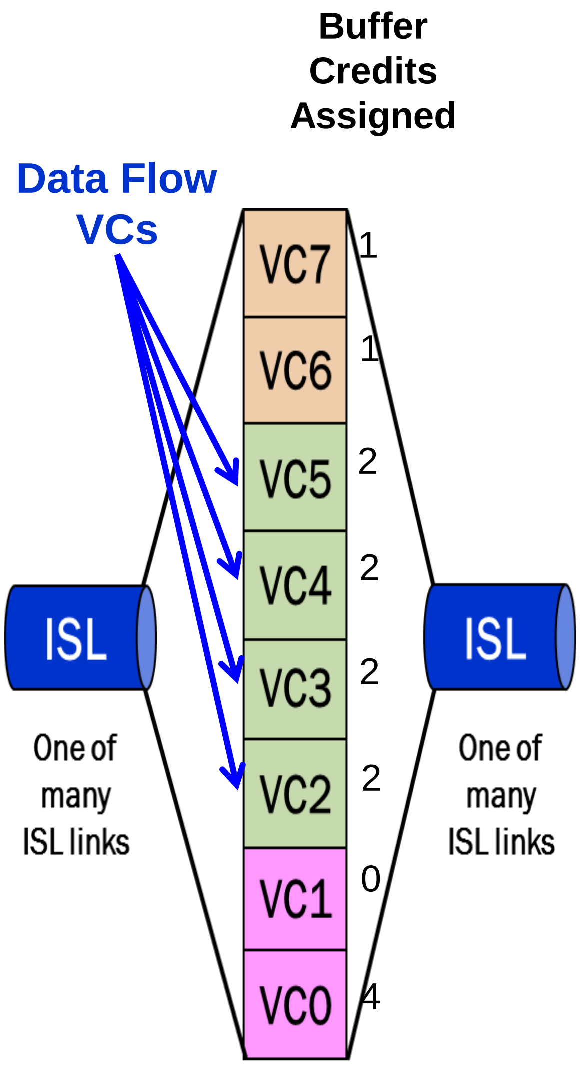

VC 2, 3, 4 and 5 can borrow a BC from the shared pool if necessary. but the limit is 3 BCs for VCs 2-5.

ISL Link Buffer Credits in Mode L0 – QoS Enabled

Physical BCs are mapped into Virtual Channel BCs

- When QoS is enabled on an E_Port, a total of 34 BCs become distributed across the VCs on each of the physical links that comprise the ISL link.
- QoS requires the Adaptive Networking License.

34 buffer credits on both of the ISL ports
- 28 BCs for the VCs
- 6 BCs shared across VCs 2-5 and 10-14
- Each data VC has 2 BCs + 1 shared BC

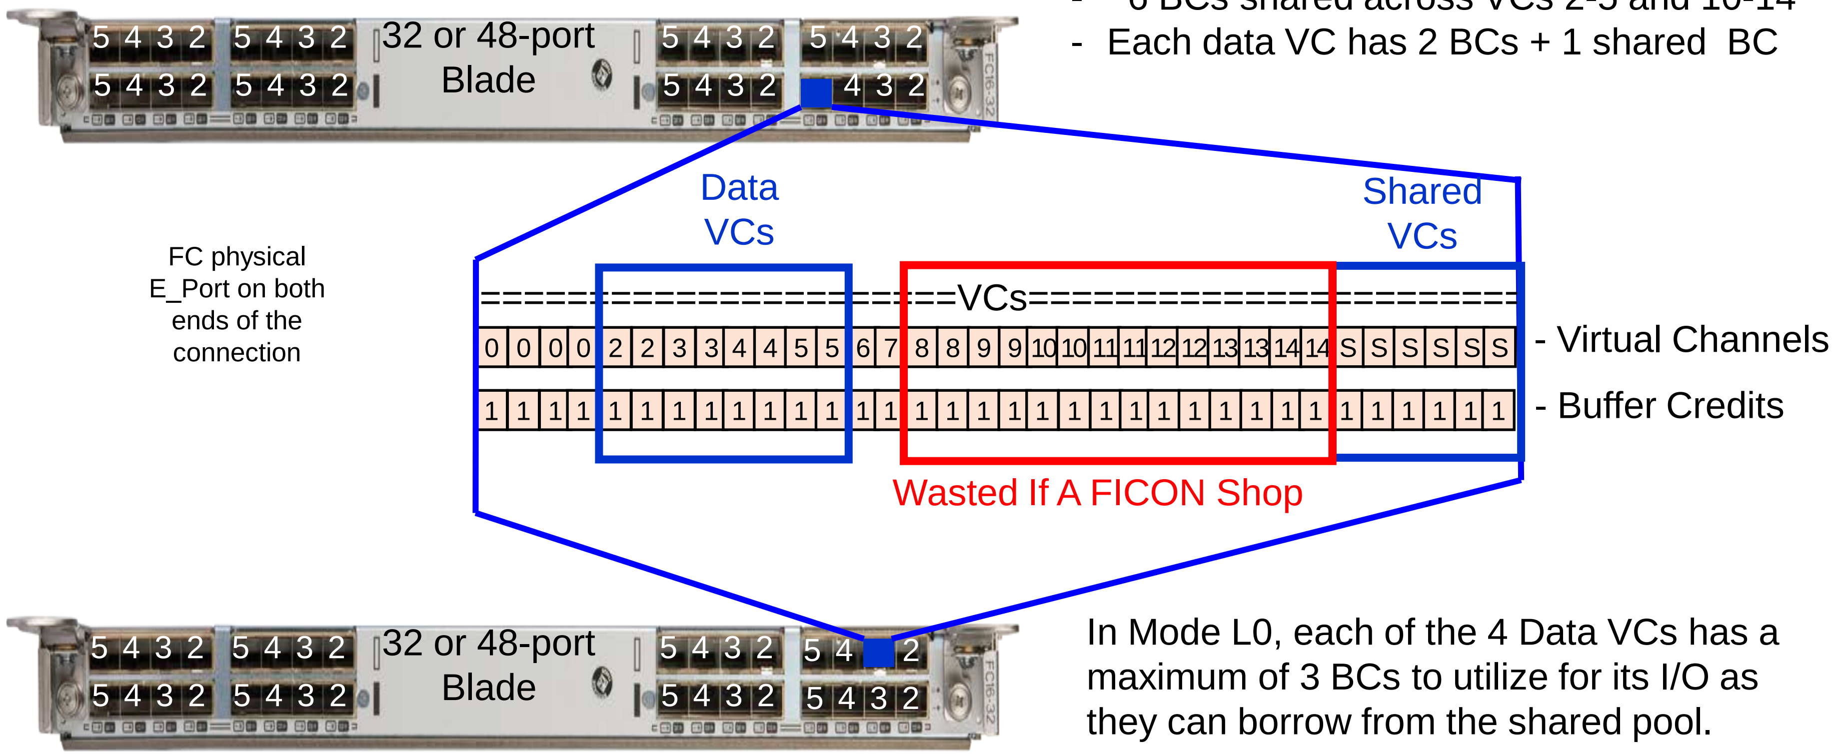

In Mode L0, each of the 4 Data VCs has a maximum of 3 BCs to utilize for its I/O as they can borrow from the shared pool.

ISL Link Buffer Credits in Mode L0 – QoS Enabled

Physical BCs are mapped into Virtual Channel BCs

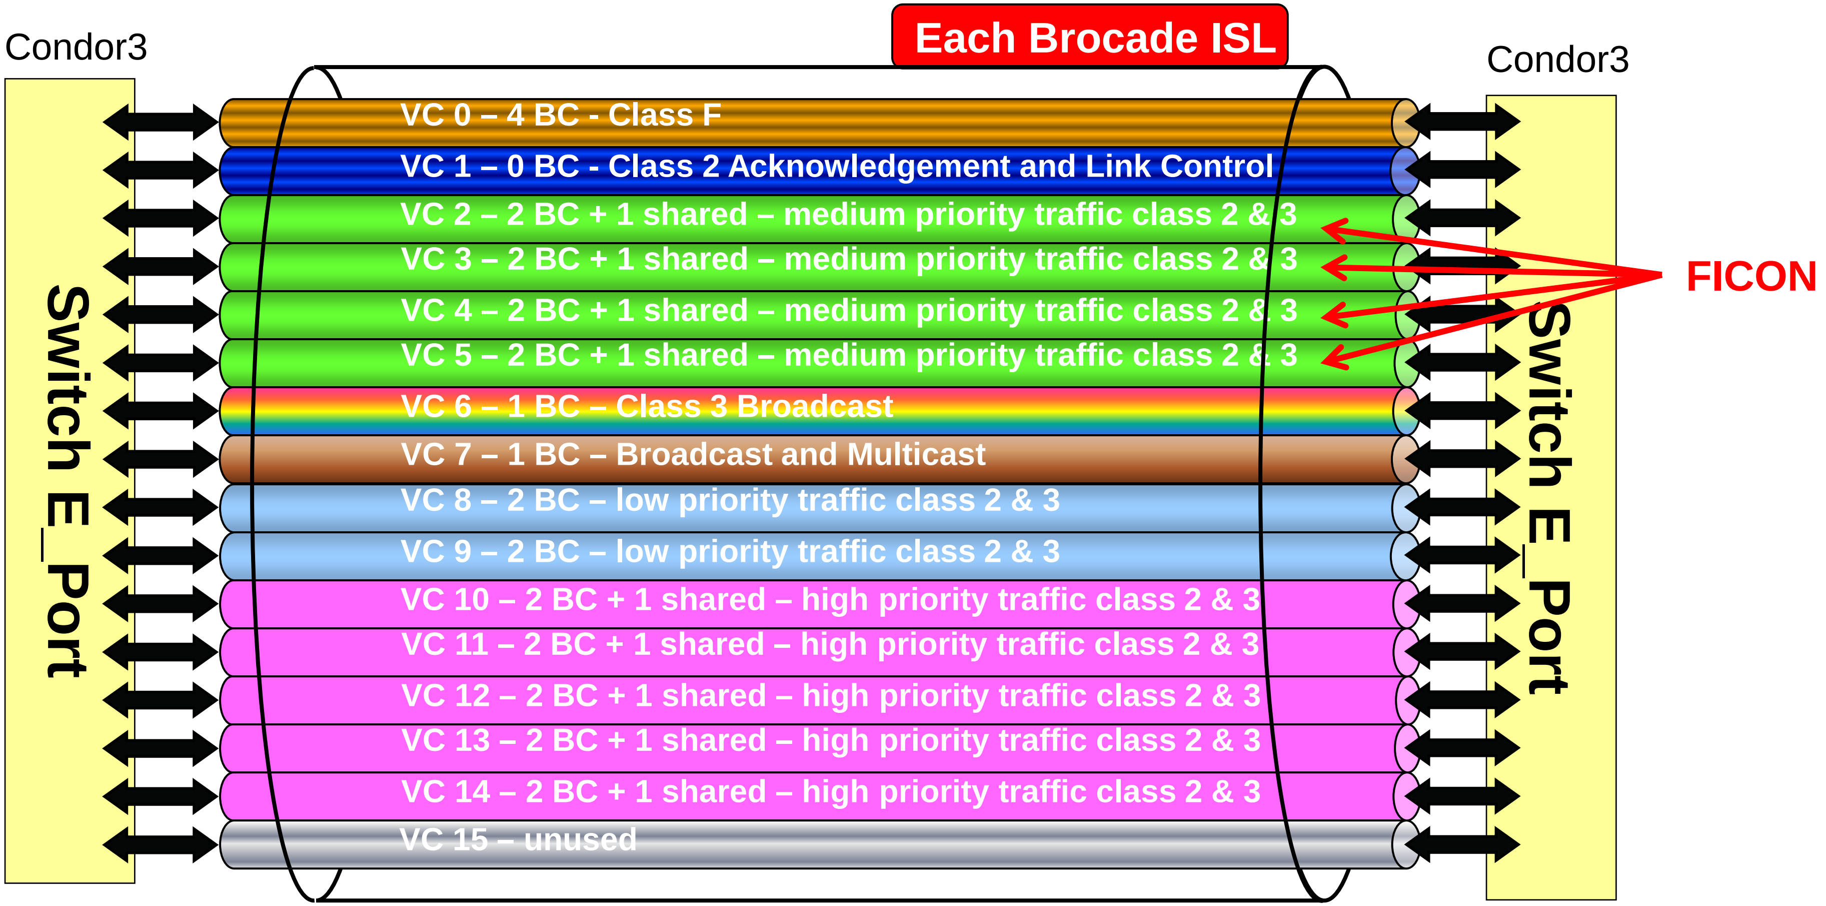

If QoS is enabled, by default, FOS provides E_Ports with 34 BCs to be used for its data I/O across 11 data VCs. The 34 BCs are divided evenly between VCs 2-5, 8-9 and 10-14.

At FOS 7.2 and later, modify this by using the portCfgEportCredits CLI command.

FOS 7.2.0b and ISL BCs in Mode L0, QoS Enabled

Physical BCs are mapped into Virtual Channel BCs

- FOS v7.2 allows buffer credit assignment even for "normal distance" (regular) E_Ports
- The new ***portCfgEportCredits*** CLI allows users to specify how many buffer credits to use.
- Prior to FOS 7.2, in L0 mode with QoS enabled, the four data VCs for FICON only had 2 buffer credits each and could borrow a BC from a shared pool.
- Now, with this new CLI command, those credits can be modified to as many as 40 BCs per VC.

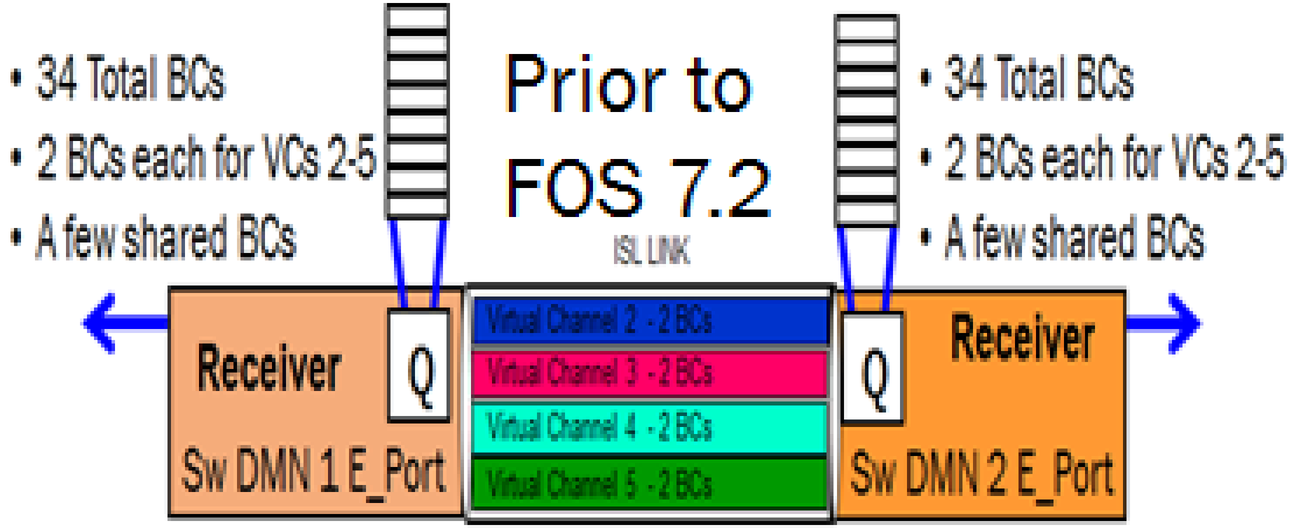

Physical ISL Link Buffer Credits in Mode LE

Extended Fabrics License Is Not Required When Using This Mode

- **Extended Mode (LE)** - LE configures the distance for an E_Port when that distance is greater than 1 km and up to 10 km. The baseline for the buffer credit calculation is one buffer credit per km at 2 Gbps. This allocation yields the following values for 10 km:
  - 40 buffer credits per port at 8 Gbps.
  - 50 buffer credits per port at 10 Gbps
  - 80 buffer credits per port at 16 Gbps

"*portcfglongdistance*" CLI Command is used to set this mode up on the ISL link pair.

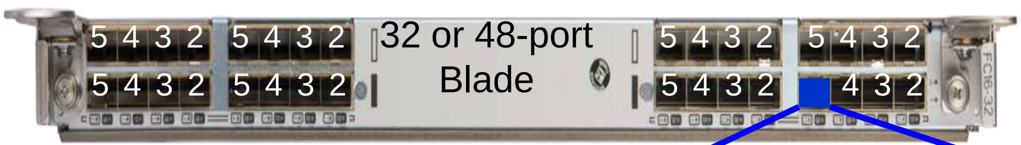

Default physical E_Port on both ends of the connection

VC:	0	1	**2**	3	4	5	6	7
BC:	4	0	**80**	0	0	0	1	1

All of the data VCs (2, 3, 4, 5) are collapsed onto VC 2

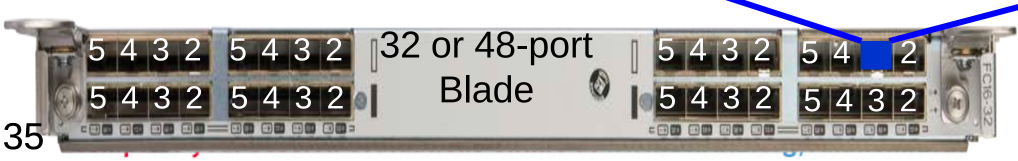

In Mode LE, the speed of the link determines how many BCs will be allocated to utilize for the ISL's I/O

Physical ISL Link Buffer Credits in Mode LD or LS
FOS 7.1.0c And the Extended Fabrics License Installed

- LD mode calculates Buffer-to-Buffer credits (BB_Credits) based on the distance measured during port initialization using a proprietary algorithm.
- LS mode originally calculates a static number of BB_Credits based solely on a user-defined distance value and framesize value.

This example assumes a distance of 16km and an average frame size of 870 bytes on a 16Gbps ISL link with compression and encryption enabled and getting 2.2:1 reduction in frame size.

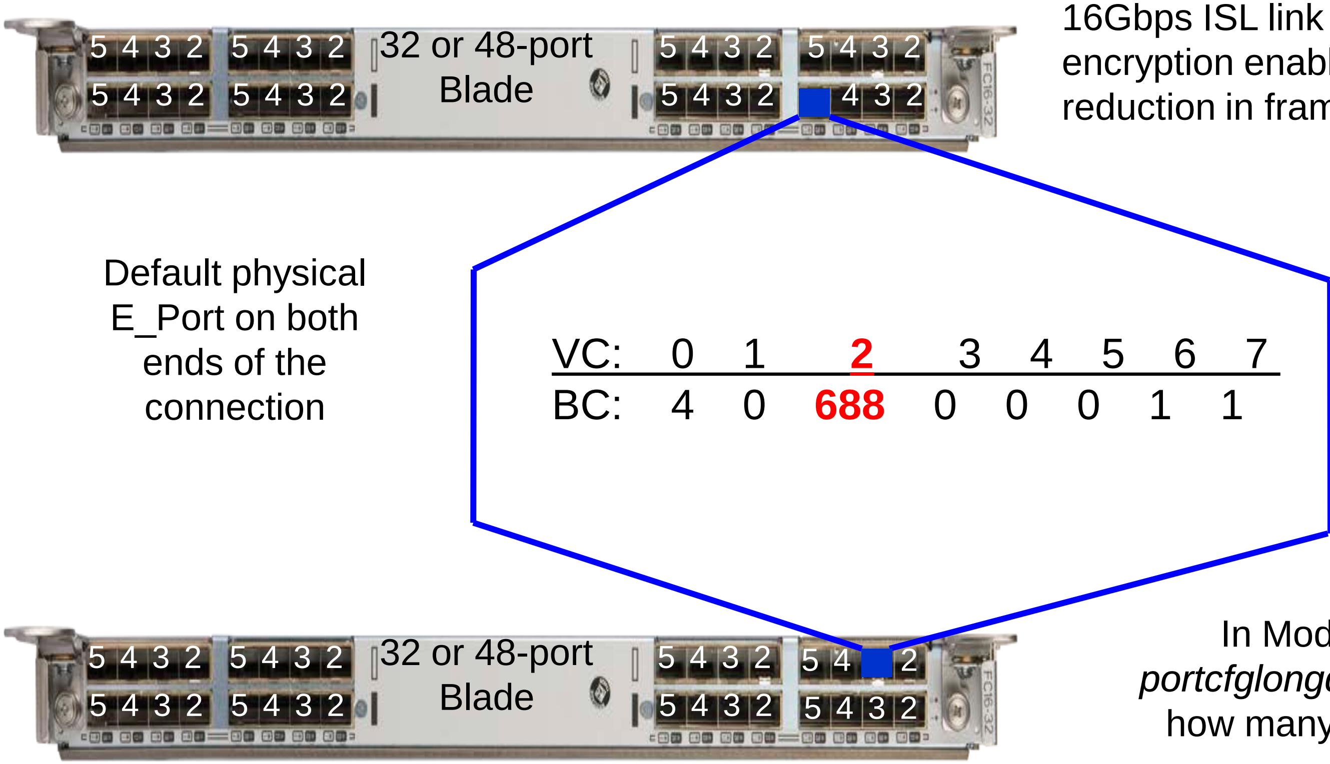

Default physical E_Port on both ends of the connection

All of the data VCs (2, 3, 4, 5) are collapsed onto VC 2

In Mode LD/LS, the CLI command *portcfglongdistance* algorithm will determine how many BCs are to be placed on VC 2

DCX Port Mapping Matrix for FICON (ALL DIR)

Using 32-port blades, no logical switches and QoS enabled

Chassis Port Number Assignment in Hex

Logical Port Layout		Slot Number																			
		1		2		3		4		5	6	7	8	9		10		11		12	
15	31	0F	1F	2F	3F	4F	5F	6F	7F	CR0	CP0	CP1	CR1	8F	9F	AF	BF	CF	DF	EF	FF
14	30	0E	1E	2E	3E	4E	5E	6E	7E					8E	9E	AE	BE	CE	DE	EE	FE
13	29	0D	1D	2D	3D	4D	5D	6D	7D					8D	9D	AD	BD	CD	DD	ED	FD
12	28	0C	1C	2C	3C	4C	5C	6C	7C					8C	9C	AC	BC	CC	DC	EC	FC
11	27	0B	1B	2B	3B	4B	5B	6B	7B					8B	9B	AB	BB	CB	DB	EB	FB
10	26	0A	1A	2A	3A	4A	5A	6A	7A					8A	9A	AA	BA	CA	DA	EA	FA
9	25	09	19	29	39	49	59	69	79					89	99	A9	B9	C9	D9	E9	F9
8	24	08	18	28	38	48	58	68	78					88	98	A8	B8	C8	D8	E8	F8
7	23	07	17	27	37	47	57	67	77					87	97	A7	B7	C7	D7	E7	F7
6	22	06	16	26	36	46	56	66	76					86	96	A6	B6	C6	D6	E6	F6
5	21	05	15	25	35	45	55	65	75					85	95	A5	B5	C5	D5	E5	F5
4	20	04	14	24	34	44	54	64	74					84	94	A4	B4	C4	D4	E4	F4
3	19	03	13	23	33	43	53	63	73					83	93	A3	B3	C3	D3	E3	F3
2	18	02	12	22	32	42	52	62	72					82	92	A2	B2	C2	D2	E2	F2
1	17	01	11	21	31	41	51	61	71					81	91	A1	B1	C1	D1	E1	F1
0	16	00	10	20	30	40	50	60	70					80	90	A0	B0	C0	D0	E0	F0

Port Numbering as 32-port blades are laid in their slots

DCX Port Mapping Matrix for FICON (ALL DIR)

Using 32-port blades, no logical switches and QoS enabled

Virtual Channel Port Number Assignment

Logical Port Layout		Slot Number																			
		1		2		3		4		5	6	7	8	9		10		11		12	
15	31	5	5	5	5	5	5	5	5	CR0	CP0	CP1	CR1	5	5	5	5	5	5	5	5
14	30	4	4	4	4	4	4	4	4					4	4	4	4	4	4	4	4
13	29	3	3	3	3	3	3	3	3					3	3	3	3	3	3	3	3
12	28	2	2	2	2	2	2	2	2					2	2	2	2	2	2	2	2
11	27	5	5	5	5	5	5	5	5					5	5	5	5	5	5	5	5
10	26	4	4	4	4	4	4	4	4					4	4	4	4	4	4	4	4
9	25	3	3	3	3	3	3	3	3					3	3	3	3	3	3	3	3
8	24	2	2	2	2	2	2	2	2					2	2	2	2	2	2	2	2
7	23	5	5	5	5	5	5	5	5					5	5	5	5	5	5	5	5
6	22	4	4	4	4	4	4	4	4					4	4	4	4	4	4	4	4
5	21	3	3	3	3	3	3	3	3					3	3	3	3	3	3	3	3
4	20	2	2	2	2	2	2	2	2					2	2	2	2	2	2	2	2
3	19	5	5	5	5	5	5	5	5					5	5	5	5	5	5	5	5
2	18	4	4	4	4	4	4	4	4					4	4	4	4	4	4	4	4
1	17	3	3	3	3	3	3	3	3					3	3	3	3	3	3	3	3
0	16	2	2	2	2	2	2	2	2					2	2	2	2	2	2	2	2

VC Allocations as 32-port blades are laid in their slots

Summary

- Cut-through frame routing, Virtual Channels and Hardware Trunking are unique Brocade features

- Combined together through best practices, these technologies automate and simplifies many of the tasks required in order to optimize the use of switched Fibre Channel and especially of storage networks using ISL links

- Whether used together or separately, these technologies provide powerful tools for maximizing efficiency, optimizing performance, increasing reliability, and reducing the cost of ownership for large FICON and open systems storage networks

- They help you build your own Fibre Channel Super Highway!

Let Me End This Session By Proudly Presenting The Industry's ONLY FICON Certification

SHARE
Technology · Connections · Results

Brocade
Certified Architect
for FICON

Mainframe-centric Classes We Teach

This FICON Certification is Unique in the Industry

Fundamentals of Brocade Mainframe Networking Seminar taught over 2 full days or 3 relaxed days

- We have been holding classes since mid-2008
- This is good for mainframers who desire to become professionally certified as FICON subject matter experts
- This seminar teaches advanced concepts and is not well suited for professionals with less than 1 year of experience

Total number of attendees at these seminars since 2008: **~550**
Total number of Brocade FICON Certifications awarded: **250+**

(There is so little FICON training available that some people attend just to get current on FICON and not for certification)

Brocade Certified Architect for FICON (BCAF)

Brocade can provide this 2 or 3 day training for YOU!

Fundamentals of Brocade Mainframe Networking Seminar

(the preparatory class for the BCAF certification)

Day 1 (8:30am to 5pm):

- Course Introduction
- Mainframe and FICON Overview
- Brocade Switching Technologies
- Design and Migration
- FICON cascading and data transmission technologies

Day 2 (8:30am to 5pm):

- Managing cascaded FICON environments
- FCIP, Data Replication and Business Continuity Networks
- FICON Implementation
- Managing and Maintaining a FICON Environment

Can be done In 3 easier days!

- Class minimum size is 8 people
- No charge if taught in the USA
- $2,500 fee for travel and living if taught overseas

- Contact David Lytle if interested: dlytle@brocade.com

Some Mainframe Social Media Sites

Please check out the Brocade Mainframe Solutions blog today:
http://community.brocade.com/t5/Mainframe-Solutions/bg-p/MainframeSolutions

Find useful information on the Brocade Fibre Channel Communities Page today:
http://community.brocade.com/t5/Fibre-Channel-SAN/bd-p/fibre

Almost 300,000 hits!

Join us in discussing issues at the Mainframe Discussion Community today:
http://community.brocade.com/t5/Mainframe-FICON/bd-p/mainframesolutionsforum

Your Eval

My Reaction!

5 = “Aw shucks. Thanks!”
4 = “Mighty kind of you!”
3 = “Glad you enjoyed this!”
2 = “A Few Good Nuggets!”
1 = “You Got a Nice Nap!”

David Lytle, BCAF

Principal Engineer

Brocade

dlytle@brocade.com

Session 14769