

CONFIDENTIAL

ENDL Letter

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

FC-LS WORKING GROUP FEBRUARY 5

After months of hiatus, the group buckled down to the ELS definition and cleanup work according to the plans laid when the FC-LS project was created.

T11.3 MEETING FEBRUARY 5 ROW

The FC-SW-3 first Public Review will close on March 15.

T11.5 MEETING FEBRUARY 5 ROW

The HBA API standard winged its way to first Public Review while the Letter Ballot plans flew fast and furiously for MIBs and FAIS.

T11 MEETING FEBRUARY 5 ROW

The annual 5-year review saw FC-10KCR, FC-FP, and FC-GS-2 withdrawn. FC-PH, FC-AL-2, FC-TAPE, and MJS were reaffirmed and 10 GFC has been published.

ADDRESSING HISTORY

Products do not always ship with the capabilities which the standard was designed to incorporate.

The first example in T11 occurred when Ancor introduced the original Fibre Channel chip. The standard included Fill Data Bytes to indicate when a frame was not an even multiple of four bytes, and it was intended to apply to any and every frame. The 2 KB frame size could be aligned easily with memory but application transfers can hardly be expected to begin on a 2K boundary.

If memory is contiguous, the Fill Data Bytes allowed a short first frame referred to as a runt to be followed by a series of full-size frames, and terminating with a runt frame. High performance systems are built with pipelining and varying sizes of cache, and there were ideas about using runts to align cache lines to fill frames.

Ancor harpooned the concept because its chip only recognized Fill Data Bytes on the last frame of a Sequence.

There was much wailing and gnashing of teeth before the members approved changing the standard to say that Fill Data Bytes 'shall only be meaningful on the last Data frame of a series of consecutive Data frames of a single Information Category within a single Sequence'.

This misstep was hardly fatal, it is not even clear that it had any effect one way or the other on the implementations or acceptance of Fibre Channel.

Some missteps are egregious, and we leave it to the reader to decide whether the following should fall into that category.

A site license is granted to subscribers so that the ENDL Letter may be copied for wide distribution within the organization.

The ENDL Letter is to be treated with the same confidentiality as internal company memos. Your co-operation is appreciated.

CONFIDENTIAL

ENDL Letter

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A site license is granted to subscribers so that the ENDL Letter may be copied for wide distribution within the organization.

The impetus for this article was sparked by mention in the Fabric Extensions Study Group that 'Fibre Channel today is stuck with just 239 switches per fabric'. The remainder of this article is not a PG-13 story, small children and those prone to crying should stop reading now. As happens in the movies, we begin with flashbacks.

Sparse addressing was never a consideration in the architecture of Fibre Channel. The fabric model was based on the telephone system of assigning addresses. This brief extract from the August 1992 Happenings report on the Fibre Channel Working Group illustrates that members relied on the familiarity of the telephone in those early days.

K.C. Chennappan (IBM) wrestled with the syntax.

"If I compare your terms to the telephone company then a Domain is the equivalent to Area Code and an Area is the Prefix on a number."

Telephone companies do not engage in sparse addressing, they assign every number they can. Sparse addressing is not even in the telecon dictionary.

In pre-internet days, most technology transfer took place at conferences, and ENDL co-sponsored a series on interfaces and storage. The Happenings report on January's 1994 Gigabit Highway Forum included explanatory material on Fibre Channel fabrics vis a vis other topologies by Kumar Malavalli. Kumar was with Hewlett Packard Canadian Network Operations at the time, and later left to become a founder at Brocade.

"Fibre Channel blends characteristics of all topologies in that it offers:

- Class 1 circuit switch*
- Class 2 and Class 3 frame switch*
- Class 4 virtual circuits for isochronous."*

When it comes to management, the requirements of switched systems are both complex and extensive. It is an area of value added and differentiation as well as a set of necessary functions. Although he did not attempt to list all the capabilities of each configuration, Kumar juxtaposed the features of one against another to show how the Fibre Channel fabric has plagiarized from prior technologies.

<i>Telephone Circuit Switched Exchange</i>	<i>Data Packet Switched Exchange</i>	<i>Fibre Channel Hybrid Exchange</i>
<i>White pages</i>	<i>Email</i>	<i>Name service</i>
<i>Yellow pages</i>	<i>Directory service</i>	<i>Directory service</i>
<i>Clock synchronize</i>	<i>Time service</i>	<i>Time service</i>
<i>Billing/accounting</i>	<i>Billing/accounting</i>	<i>Management service</i>
<i>Network management</i>	<i>Network management</i>	<i>Network management</i>

The ENDL Letter is to be treated with the same confidentiality as internal company memos. Your co-operation is appreciated.

CONFIDENTIAL

ENDL Letter

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A site license is granted to subscribers so that the ENDL Letter may be copied for wide distribution within the organization.

Features	Features	Features
- Hunt group	- Datagram	- Intermix
- Call waiting	- Multicast	- Multicast
- Call parking	- Broadcast	- Hunt group
- Call forwarding	- Virtual circuit	- Virtual circuit
- Conference call	- Closed user group	- Closed user group
- Alternate routing	- Alternate routing	- Stacked connection

Kumar focused on the architectural overview and structures of the fabric, and it was left to cohort Bent Stoevase (Hewlett Packard Canadian Network Operations) to 'describe how a fabric is made up of modular building blocks named Fabric Elements'.

A piece of magic handled by the fabric is port level addressing. Assignments of the high-order 16 bits are made automatically by the fabric, and the last 8 bits of the 24-bit address are left to be locally administered.

"U.S. and Canada use a 10-position address of Area Code, Prefix and Local. The 24-bit Fibre Channel addresses are similar to the AAA-PPP-LLLL of the telephone system."

- Up to 240 Domains 8-bits DDDDDDD-*
- Up to 240 Areas 8-bits DDDDDDD-AAAAAAA-*
- Up to 256 Local 8-bits DDDDDDD-AAAAAAA-LLLLLLL*

"Just like 411 and 911 and other country-wide numbers, there have to be addresses for well-known services, and many are available. There are 4096 addresses for fabric-assisted functions in a Domain."

"The really hard part of fabrics from the outside looking in is routing."

"It is no easier on the inside. All the decisions on how to get frames from one port to another have to happen at speeds of up to 100 MBs."

"Latency is unacceptable unless the fabric wormholes at each switch point so delays cannot build up to the point of affecting system performance."

- o Routing paths are made automatically by the fabric*
- o The internal structure of remote Areas is invisible*
- o The internal structure of remote Domains is invisible*
- o Routing paths are established to ports within an Area*
- o Routing paths are established to other Areas within a Domain*
- o Routing paths are established to other Domains within the fabric*
- o All the Inter-Element Links joining Areas are used for routing*
- o All the Inter-Element Links joining Domains are used for routing*
- o Routing path selection may be modified/overridden by network management*



Take note here of key words like 'internal structure....is invisible'.

The ENDL Letter is to be treated with the same confidentiality as internal company memos. Your co-operation is appreciated.