VCAT/LCAS IN A CLAMSHELL

GREG BERNSTEIN, DIEGO CAVIGLIA, RICHARD RABBAT, AND HUUB VAN HELVOORT

INTRODUCTION

Virtual concatenation (VCAT) is a standardized layer 1 inverse multiplexing technique that can be applied to the optical transport network (OTN) [1], synchronous optical network (SONET) [2], synchronous digital hierarchy (SDH) [3], and plesiochronous digital hierarchy (PDH) [4] component signals. By inverse multiplexing — sometimes referred to as concatenation — we mean a method that combines multiple links at a particular layer into an aggregate link to achieve a commensurate increase in available bandwidth on that aggregate link. More formally, VCAT essentially combines the payload bandwidth of multiple path layer network signals (or trails) to support a single client (e.g., Ethernet) layer link. Other well-known standardized inverse multiplexing techniques include Multi Link PPP [5] and Ethernet's Link Aggregation mechanism as documented in chapter 43 of [6].

While any inverse multiplexing scheme is about "more bandwidth," VCAT/LCAS is a general technique that can enable a fairly broad range of network features such as:

Right sizing bandwidth for data applications: Circuit-switched multiplex hierarchies and most link technologies are fairly inflexible in terms of the bandwidth increments they offer. For example, in the SDH hierarchy we have a VC-3 at approximately 50 Mb/s or a VC-4 at 150 Mb/s. For the carriage of a full-rate 100 Mb/s Ethernet connection, the VC-3 is 50 percent too small, while the VC-4 is 50 percent too large. VCAT provides just the "right size" pipe for this application: a VC-3-2v.

Extracting bandwidth from a mesh network: Given an end-to-end bandwidth demand between a source and a destination, and a mesh network topology, there may be enough total bandwidth across the network to meet the demand, but not along a single route. VCAT allows us to "extract" the required bandwidth from a mesh since it can "glue" together pipes that follow different paths through the network to give a larger pipe that meets the requested demand.

Bandwidth on demand and IP traffic engineering: The Link Capacity Adjustment Scheme (LCAS) companion to VCAT allows for hitless resizing of bandwidth between two circuit endpoints. Probably the most common method of IP layer traffic engineering involves adjustment of link weights in a link state routing protocol. The problem with such techniques is that they are disruptive to the network (routing protocols must converge, and all route tables need to be recalculated); hence, such optimization is generally done on long timescales such as weeks and months. One of the main differences between VCAT and the other mentioned inverse multiplexing standards is that VCAT works at layer 1 rather than at the data link layer; that is, VCAT works with "circuits" and the others with layer 2 packets. Changes to bandwidth via VCAT/LCAS between routers will not alter IP layer topology; hence, with VCAT/LCAS we can respond to shorter timescale optimizations on a per IP link basis [7].

Painless regrooming: When connections need to be rerouted due to maintenance or to make efficient use of network resources the process, known as regrooming, generally impacts user traffic. Although proprietary "make before you break" schemes exist, VCAT/LCAS enables a hitless method for regrooming by first adding additional components that have been set up on the new desired path, then first removing the old components from the VCAT group and then releasing the resources from the network.

New forms of protection/restoration and graceful degradation: VCAT/LCAS itself provides a graceful degradation (reduction of bandwidth) in response to VCAT group component failures. But additional techniques have been developed [8] that allow more flexibility than existing protection/restoration schemes in trading off between network bandwidth efficiency, restoration time, and robustness to failure scenarios [9].

VCAT performs inverse multiplexing by octet/byte deinterleaving of the encapsulated client bitstream. As such it operates below the packet/frame level. Each frame/packet will therefore "travel" over all members of the VCAT group, and a fault in any of the members of the VC-n-Xv hits every Xth byte in each packet/frame. With LCAS enabled the failed member is temporarily taken out of the service providing set of the VCAT group, until the fault is repaired. Due to this octet/byte deinterleaving, VCAT introduces an insignificant processing delay into the transmission path. The propagation time for the aggregate signal will correspond to that of the longest component signal.

Figure 1 illustrates how incoming client traffic, in this case an Ethernet frame, is transported via VCAT in a transport network. The incoming Ethernet frame (for the sake of simplicity only six bytes of the frame are depicted) is inverse-multiplexed by VCAT into three different VCAT members. In Figure 1 the incoming Ethernet frame is spread across the three VCAT members, that is, bytes 1 and 4 are carried by VCAT member number 1, bytes 2 and 5 by member number 2 and bytes 3 and 6 by member number 3. In a failure of VCAT member 2, bytes 2 and 5 are lost; thus, it is not possible to rebuild the original incoming Ethernet frame.

VCAT SIGNALS, CAPABILITIES, AND LIMITATIONS

SDH/SONET VCAT SIGNALS AND COMPONENTS

In SDH (and similarly in SONET) VCAT can be applied to the following component time division multiplex (TDM) signals referred to as Virtual Containers (VCs) (and not to be confused with virtual circuits): VC-11, VC-12, VC-2, VC-3, and VC-4.

Note that when reading the VCAT and LCAS references the term "frame" is generally used to describe the repetitive structure of TDM signals and not to describe a layer 2 packet. To simplify high-speed hardware aggregation of these signals, only like component signals can be aggregated into a VCAT group. The aggregate signals are named and characterized in Table 2 extended from [3, Table 11-4].

Since VCAT is an inverse multiplexing technique, intermediate SONET/SDH transport network nodes do not need to support these VCAT signals explicitly since it is the job of the VCAT end systems to reassemble the aggregate signal. The only requirement on the SONET/SDH network is to be able to transport the individual component signals of Table 1.

PDH VCAT SIGNALS AND COMPONENTS

VCAT can be applied to the following PDH signals as specified in [4]: DS1, E1, E3, and DS3. Similar to the SONET/SDH case, these component signals can only be combined with like signals to produce aggregates. The PDH VCAT groups use a similar notation to the SDH VCAT signals by using the commonly used designations shown in Table 2.

 0163-6804/06/$20.00 © 2006 IEEE

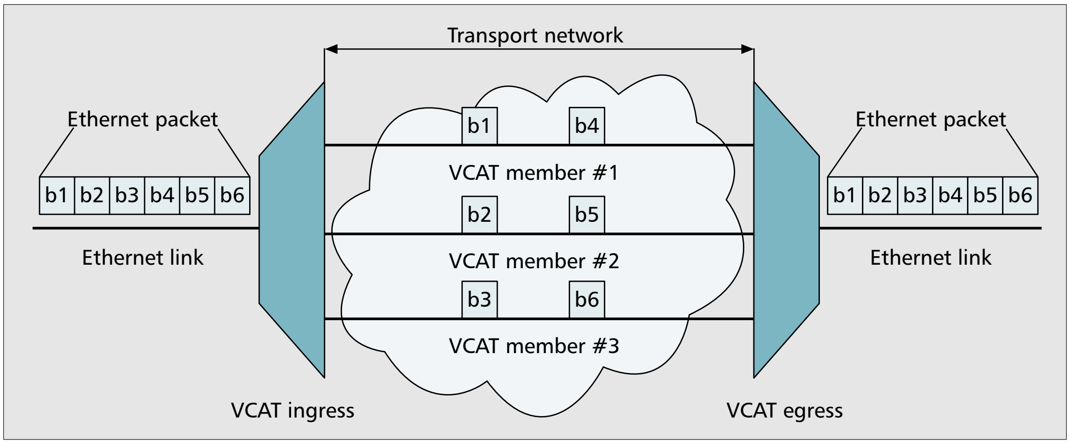

FIGURE 1. VCAT inverse multiplexing.

OTN VCAT SIGNALS AND COMPONENTS

Concatenation in the OTN is realized by means of VCAT of optical channel payload unit (OPU) signals. OPU*k* signals (*k* = 1, 2, 3) can be concatenated into OPU*k*-*X*v aggregates. The aggregate signals are named and characterized in Table 3 (adapted from Table 6-3, G.8012).

Note that the last row in Table 3 is not a misprint. Reference [1] does indeed permit VCAT of up to 256, 40 Gb/s, ODU3 signals to produce an aggregate link, ODU3-256v, with a capacity of over 10 Tb/s! At the time of this writing the authors do not currently know of any actual implementations, but it should be noted that the standard appears quite "future-proof."

VCAT CAPABILITIES AND LIMITATIONS

With any inverse multiplexing technique two important issues come up: how to prevent packet reordering, and delay compensation limits. For example, Ethernet's link aggregation scheme prevents reordering by restricting "conversations" to a single link. This means that the total aggregate bandwidth is not available to a single flow. MLPPP and VCAT prevent reordering in a way that imposes no limits on the bandwidth delivered to a single flow. Since VCAT works with circuits it does not have to deal with queuing induced differential delays between components. In fact, since most circuit-switched technologies have very low switching latency, most differential delays experienced by VCAT component signals are due to propagation. The maximum differential delay that can be accommodated by the standards is given in Table 4. Actual implementation can choose to provide much less differential delay compensation and frequently do so to save on memory requirements.

As mentioned in [9], the ability to compensate for over 200 ms of differential delay compares favorably with the circumference of the Earth and some rather paranoid disjoint paths.

THE LCAS PROTOCOL

The Link Capacity Adjustment Scheme for VCAT signals is a protocol for dynamically and hitlessly changing (i.e., increasing and decreasing) the capacity of a VCAT group. LCAS also provides survivability capabilities, automatically decreasing the capacity if a member of the VCAT group experiences a failure in the network, and increasing the capacity when the network fault is repaired. LCAS itself provides a mechanism for interworking between LCAS and non-LCAS VCAT endpoints. VCAT does not require LCAS for its operation.

We find analogous mechanisms in other inverse multiplexing technology such as the Link Control Protocol (LCP) used in MLPPP [5] and the Link Aggregation Control Protocol (LACP) used in Ethernet link aggregation [6]. It needs to be emphasized that none of these mechanisms are responsible for establishing the component links. Indeed, these protocols run over the component links themselves. Hence, LCAS functionality does not overlap or conflict with generalized multiprotocol label switching's (GMPLS') routing or signaling functionality for the establishment of component links or entire VCAT groups. LCAS instead is used to control whether

SDH VCAT type	Component Signal	X range	Capacity (kb/s)
VC-11-Xv	VC-11	1 to 64	1600 to 102 400
VC-12-Xv	VC-12	1 to 64	2176 to 139 264
VC-2-Xv	VC-2	1 to 64	6784 to 434 176
VC-3-Xv	VC-3	1 to 256	48 348 to 12.5 Gb/s
VC-4-Xv	VC-4	1 to 256	149 760 to 38.3 Gb/s

TABLE 1. SDH VCAT signals.

PDH VCAT type	Component signal	X range	Capacity (kb/s
DS1-Xv	DS1	1 to 16	1533 to 24 528
E1-Xv	E1	1 to 16	1980 to 31 680
E3-Xv	E3	1 to 8	33 856 to 270 848
DS3-Xv	DS3	1 to 8	44 134 to 353 072

TABLE 2. PDH VCAT signals.

OTN VCAT type	Component signal	*X* range	Capacity (kb/s)
OPU1-Xv	OPU1	1 to 256	2,488,320 to 637,009,920
OPU2-Xv	OPU2	1 to 256	~9,995,277 to ~2,558,709,902
OPU3-Xv	OPU3	1 to 256	~40,150,519 to ~10,278,532,946

TABLE 3. OTN component and VCAT signals.

a particular component signal is actually put into service carrying traffic for the VCAT group.

Although we are used to PDH and SONET/SDH signals being bidirectional, LCAS actually works on unidirectional components in a VCAT group with the proviso that there is at least one return component for conveyance of LCAS messages. The forward and return signal capacities are allowed to be different. As viewed from LCAS' point of view, the source end of each component can have the following states:

- IDLE state — This member is not provisioned to participate in the concatenated group.
- NORM state — This member is provisioned to participate in the concatenated group and has a good path to the sink end.
- DNU state — (Do Not Use) This member is provisioned to participate in the concatenated group and has a failed path to the sink end.
- ADD state — This member is in the process of being added to the concatenated group.
- REMOVE state — This member is in the process of being deleted from the concatenated group.

LCAS provides for graceful degradation of failed links by having the sink end report back the receive status of all member components. In the case of a reported member failure, the source end will stop using the component and send an LCAS control word to the sink end that it is not transmitting data on that component. The worst case notification times, not including propagation delays, for the different VCAT signals discussed here are given in Table 4. These values were obtained from [1, 3], and derived from information in [4].

CONCLUSION AND NEXT STEPS

We have given a quick overview of VCAT/LCAS technology and just a few examples of its applications. Work on enhancing GMPLS/G.ASON has recently been undertaken at the Internet Engineering Task Force (IETF) [10]. From [10] a VCAT/LCAS-friendly control plane would include:

- Discovery of VCAT: VCAT sources can only communicate with VCAT-capable sinks. Hence, the VCAT capabilities of PDH, SDH, or OTN path termination points need to be known.
- Discovery of LCAS: LCAS offers additional functionality between VCAT capable sources and sinks. Hence, the LCAS capabilities of VCAT-enabled path termination points can be useful to know in advance of component signal setup.
- VCAT group identification: Since we can have more than one VCAT group per GMPLS link, there is currently an irresolvable ambiguity about when disjoint member connections are set up or dynamic resizing is applied.

With these enhancements to GMPLS, the potential of the exciting new technology of VCAT/LCAS should come closer to full realization.

REFERENCES

[1] ITU-T Rec. G.709, "Interfaces for the Optical Transport Network (OTN)," Mar. 2003.
[2] ANSI T1.105-2001, "Synchronous Optical Network (SONET) — Basic Description Including Multiplex Structure, Rates, and Formats," 2001.
[3] ITU-T Rec. G.707, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)," Dec. 2003.
[4] ITU-T Rec. G.7043, "Virtual Concatenation of Plesiochronous Digital Hierarchy (PDH) Signals," July 2004.
[5] Sklower *et al.*, "The PPP Multilink Protocol (MP)," RFC 1990, Aug. 1996.
[6] IEEE Std. 802.3, "Information Technology — Telecommunications and Information Exchange between Systems — Local and Metropolitan Area Networks — Specific Requirements — Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," Mar. 2002.
[7] G. Bernstein and L. Ong, "IP Bandwidth on Demand via Optical Networks," to be published, http://www.grotto-networking.com/pages/IP_BoD.pdf
[8] C. Ou *et al.*, "Survivable Virtual Concatenation for Data over SONET/SDH in Optical Transport Networks," to appear, *IEEE/ACM Trans. Networking*.
[9] G. Bernstein, B. Rajagopalan, and D. Saha, *Optical Network Control: Architecture, Protocols and Standards*, Addison-Wesley, 2004.
[10] G. Bernstein, D. Caviglia, and R. Rabbat, "Operating Virtual Concatenation (VCAT) and the Link Capacity Adjustment Scheme (LCAS) with Generalized Multi-Protocol Label Switching (GMPLS)," Internet draft, draft-bernstein-ccamp-gmpls-vcat-lcas-01, Oct. 2005.

ADDITIONAL READING

[1] ITU-T Rec. G.7042, "Link Capacity Adjustment Scheme (LCAS) for Virtual Concatenated Signals," Feb. 2004.
[2] H. van Helvoort, *Next generation SDH/SONET: Evolution or Revolution?*, Wiley, 2005.

BIOGRAPHIES

GREG BERNSTEIN is with Grotto Networking.

DIEGO CAVIGLIA is with Marconi.

RICHARD RABBAT is with Fujitsu.

HUUB VAN HELVOORT is a networking consultant.

VCAT signal	Max diff. delay	LCAS notification time
VC-11-*X*v, VC-12-*X*v, VC-2-*X*v	256 ms	128 ms
VC-3-*X*v, VC-4-*X*v	256 ms	64 ms
DS1-*X*v	384 ms	96 ms
E1-*X*v	256 ms	64 ms
E3-*X*v	255 ms	2 ms
DS3-*X*v	217 ms	1.7024 ms
ODU1-*X*v	411 s	1.567 µs
ODU2-*X*v	102 s	390 µs
ODU3-*X*v	25.4 s	97 µs

TABLE 4. Differential delay limits and LCAS notification times for the various VCAT signals.