Click for more about Monitoring IBM Enterprise Extender: www.sdsusa.com/vip/ee

Staying On Top Of
Enterprise Extender
Proactive Monitoring, Management & Optimization

Software Diversified Services (SDS)

Perceptive Management of HPR & EE with SDS VIP

Real-time, detailed data on end-to-end EE path configurations, connection availability, congestion control, path swaps & much more …

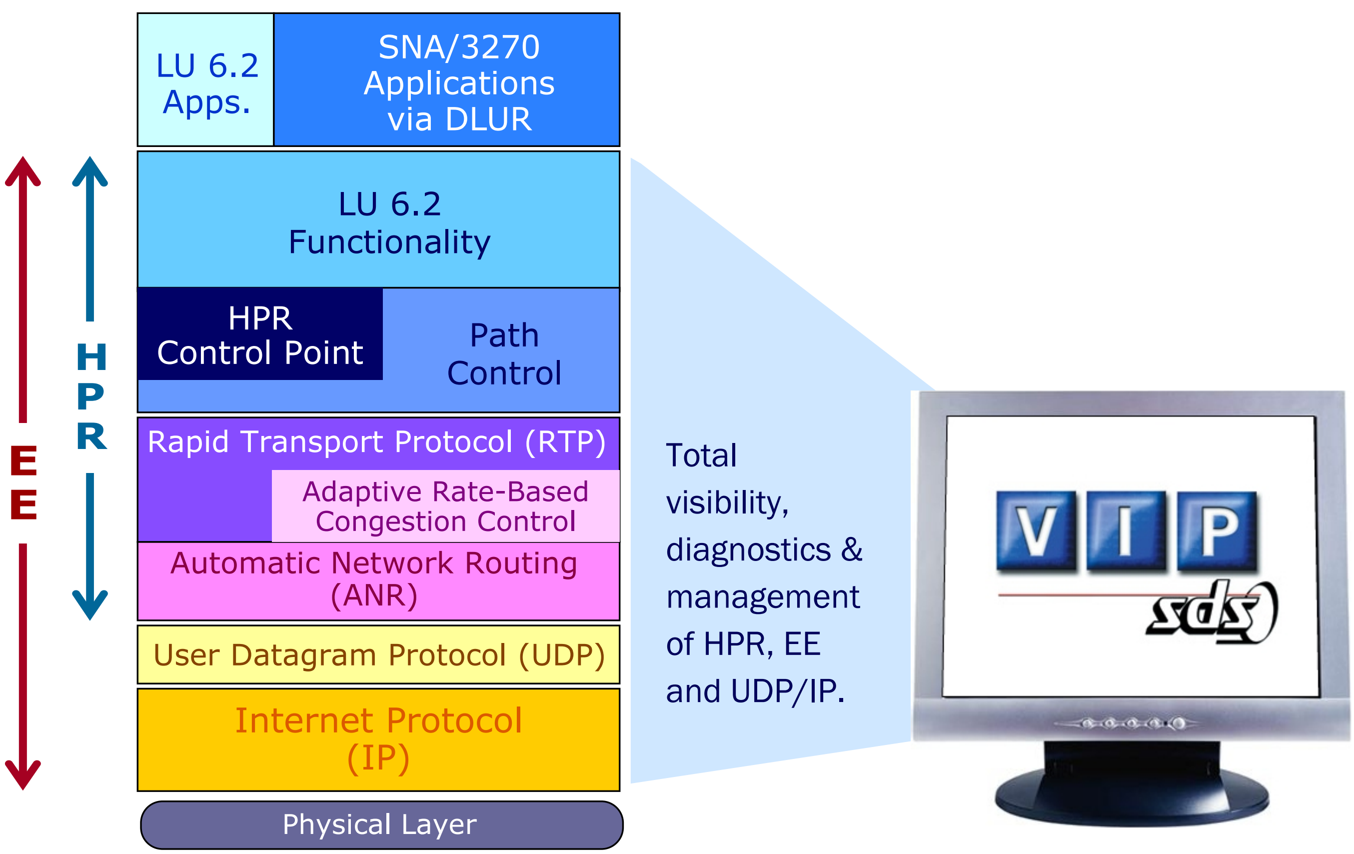

Developed for SDS
by Anura ['SNA'] Gurugé
January 2009

SOFTWARE DIVERSIFIED SERVICES

www.sdsusa.com

Table of Contents

Preamble: Overview of EE and HPR 3

VIP's Specialized Support for EE and HPR 5

Function Overlaps Necessitate Specialized Monitoring 8

Resolving Real-Life EE Problems with VIP 10

VIP's Agent/Server 'Green' Architecture 16

Partial Repertoire of VIP's Network Monitoring & Management Capabilities 17

The Bottom Line 19

Selected Glossary 20

Software Diversified Services 21

The Author 21

Key Acronyms

ARB	Adaptive Rate-Based (Congestion Control)
ANR	Automatic Network Routing
APPN	Advanced Peer-to-Peer Networking
COS	Class of Service
DLSw	Data Link Switching
DLUR	Dependent LU Requestor
EE	Enterprise Extender
HPR	High Performance Routing
IP	Internet Protocol
LU	Logical Unit
NMI	Network Management API
RTP	Rapid Transport Protocol
SNA	Systems Network Architecture
UDP	User Datagram Protocol
VIP	Vital Signs VisionNet IP Monitor
VTAM	Virtual Telecommunications Access Method

Staying On Top Of Enterprise Extender

Proactive Monitoring, Management & Optimization

Software Diversified Services

Perceptive Management of HPR & EE with SDS VIP

Trying to monitor EE is like observing a swan glide across a lake. All of the activity taking place below the surface may not be apparent.

EE, codified in RFC 2353 as *'APPN/HPR in IP networks'* in 1998, was synthesized by a multi-company committee headed by IBM. To think of it merely as 'SNA-over-IP,' or even 'APPN-over-IP,' is misleading, to say the least. DLSw and tn3270(E) represent 'SNA-over-IP.' EE, on the other hand, is very much 'HPR-over-UDP/IP.' And that from a monitoring and management perspective is the rub. To stay on top of EE networking you need more than just traditional SNA/APPN- and TCP/IP-oriented tools and expertise.

HPR is to SNA what a Segway is to a bicycle. They will both get you from A to B, but one is considerably more nimble, self-regulating and high-tech. HPR, developed in the early 1990s, was IBM's last-ditch attempt to make SNA/APPN more competitive with the then fast-prevailing IP. HPR at long last provided SNA with dynamic, alternate-routing [albeit still not by intermediate nodes]. It also included a powerful, feedback-driven, adaptive rate-based (ARB) congestion control mechanism and a label-based, Layer 2, routing scheme similar to that used by Token-Ring LAN source-route bridging (SRB). [See Fig. 1 on page 4.]

HPR totally supplanted SNA/APPN's traditional 'intermediate session routing' functionality with three new, function-rich components – as depicted in Fig. 1. These three HPR routing components, in the context of EE, then rely upon UDP, a connectionless protocol often blocked by firewalls, for their end-to-end network transport. UDP, in its turn, uses IP, with all of its networking magic, as its underlying communications mechanism. Suffice to say that EE networking, with HPR, UDP and IP all in play, is sophisticated and thus complex – if not convoluted.

Thus, in order to monitor, manage and optimize EE networking, you have to have specialized, EE-specific tools, such as SDS' VIP, that understands and supports the HPR protocols as well as their interplay with UDP/IP. Moreover, given its highly dynamic nature, the management data pertaining to EE has to be extracted, collated and analyzed quickly and efficiently – *in real-time*. Otherwise, you will invariably be working with inaccurate, elapsed data that does not reflect the true, current status of the EE network. VIP exploits IBM's strategic and recommended **Network Management API (NMI)** to obtain extensive EE data directly from VTAM, IP stacks and z/OS.

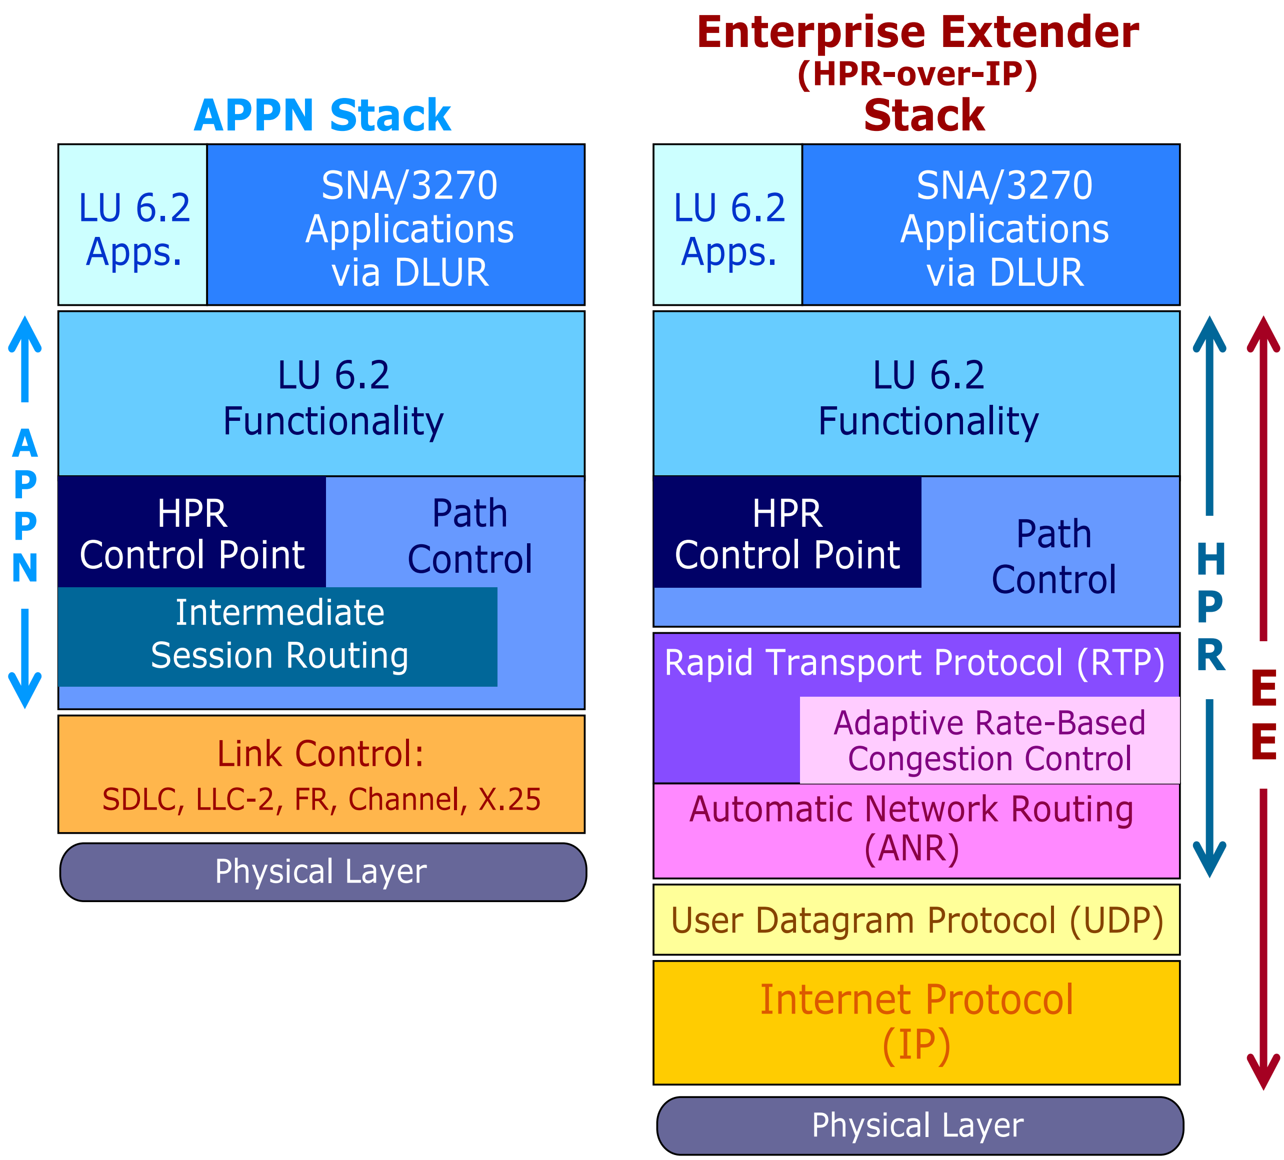

Figure 1: Enterprise Extender's protocol stack consists of a potent cocktail of protocols – hence the need for EE-specific management tools such as those included in SDS' VIP monitor.

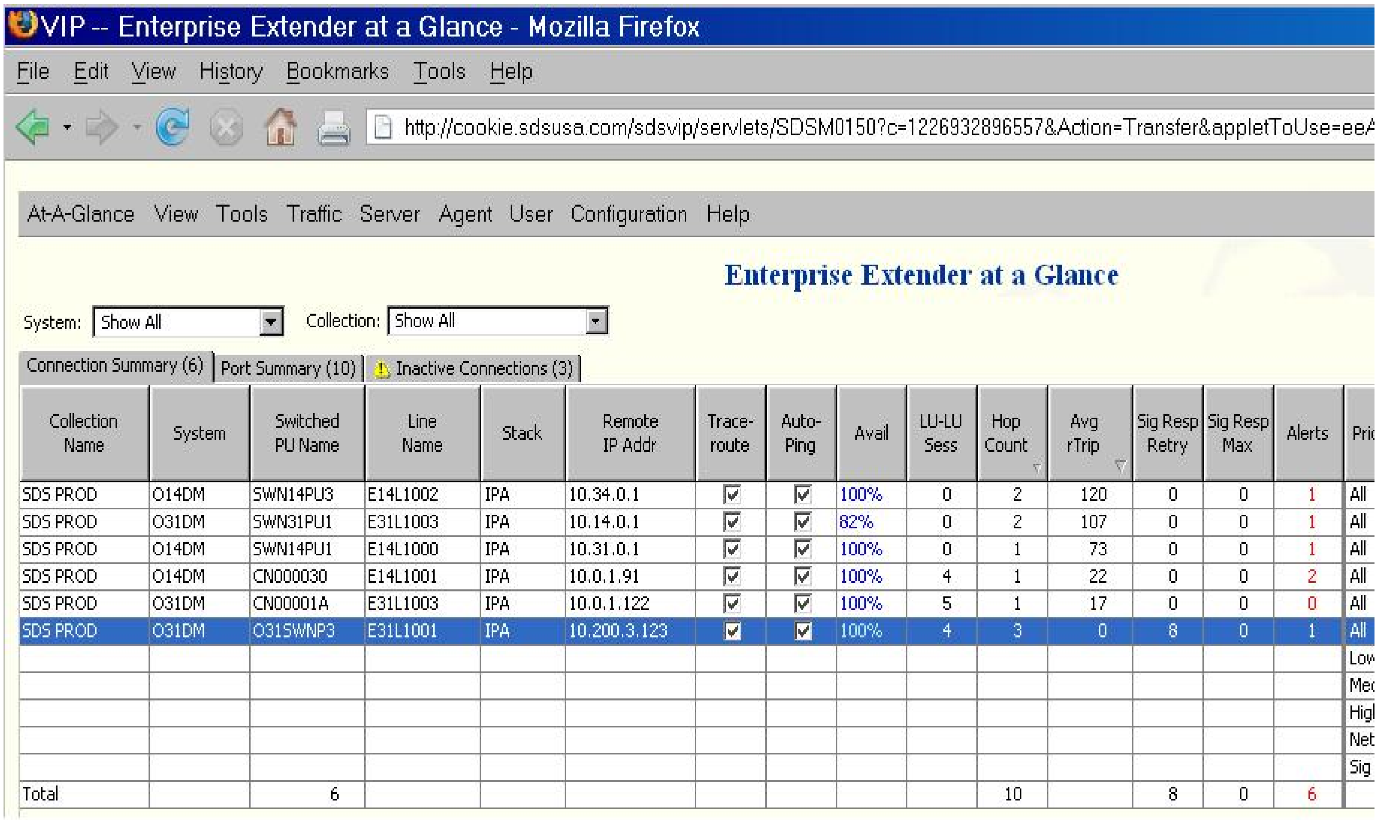

Collection Name	System	Switched PU Name	Line Name	Stack	Remote IP Addr	Trace-route	Auto-Ping	Avail	LU-LU Sess	Hop Count	Avg rTrip	Sig Resp Retry	Sig Resp Max	Alerts	Pri
SDS PROD	O14DM	SWN14PU3	E14L1002	IPA	10.34.0.1	☑	☑	100%	0	2	120	0	0	1	All
SDS PROD	O31DM	SWN31PU1	E31L1003	IPA	10.14.0.1	☑	☑	82%	0	2	107	0	0	1	All
SDS PROD	O14DM	SWN14PU1	E14L1000	IPA	10.31.0.1	☑	☑	100%	0	1	73	0	0	1	All
SDS PROD	O14DM	CN000030	E14L1001	IPA	10.0.1.91	☑	☑	100%	4	1	22	0	0	2	All
SDS PROD	O31DM	CN00001A	E31L1003	IPA	10.0.1.122	☑	☑	100%	5	1	17	0	0	0	All
SDS PROD	O31DM	O31SWNP3	E31L1001	IPA	10.200.3.123	☑	☑	100%	4	3	0	8	0	1	All
															Low
															Med
															High
															Net
															Sig
Total		6								10		8	0	6	

Figure 2: A *portion of* VIP's information-rich 'EE at a Glance' screen, in 'Port Summary' mode, providing an instant summary of network availability and status – with extensive drill-down options.

VIP's Specialized Support for EE and HPR

SDS VIP provides detailed, *end-to-end visibility* of VTAM-centric EE networks combined with an X-ray like, piercing view of the HPR-related activity taking place within a *VTAM node*. Thus, from an HPR and EE standpoint, VIP diligently monitors not just network activity but also the workings of the HPR stack within a VTAM instance. This dual-mode, VTAM and network monitoring done by VIP is key because an accurate, *real-time reading* of the workings of the HPR protocols, in particular ARB-based congestion control, is only possible by watching the changes in parameters at the VTAM HPR node. Hence the futility of trying to realize business-critical monitoring of EE with a tool, such as a packet tracing product, that only sees network traffic and has no cognizance, whatsoever, of what the VTAM HPR stack is doing.

SDS VIP's industry-leading repertoire of EE and HPR management capabilities includes:

- ►'At a glance' status/availability overviews of **EE connections** and ports [Fig. 2], with any pending alerts highlighted, with extensive 'drill down' options, including on-the-fly, instant-result 'traceroute' [Fig. 4] and 'auto-ping' features, that can be freely invoked with a right-click of the mouse.

► 'At a glance' review of an array of **HPR data**, per active connection, including *color-coded* indication of ARB congestion status, number of nodes involved (i.e., hop count), number of path switches (i.e., HPR reroutes) that have occurred, active LU-LU sessions, current data send rate (in KB/sec) and outstanding alerts [Fig. 3]. This screen, which requires horizontal scrolling in order to view all of the available columns, also provides precise information about network IDs, remote PU names, extent of ARB congestion control, packet send/receive rates (with retransmission and in-queue counts), pending acknowledgements, COS designation and transmission group (TG) availability. The level of HPR data made available by VIP is unprecedented – as is the easy-to-follow presentation, in particular the color-coded congestion indication. The VIP 'HPR at a Glance' screen will satisfy the monitoring and management needs of those new to EE as well as those that are experts – not to mention those in-between.

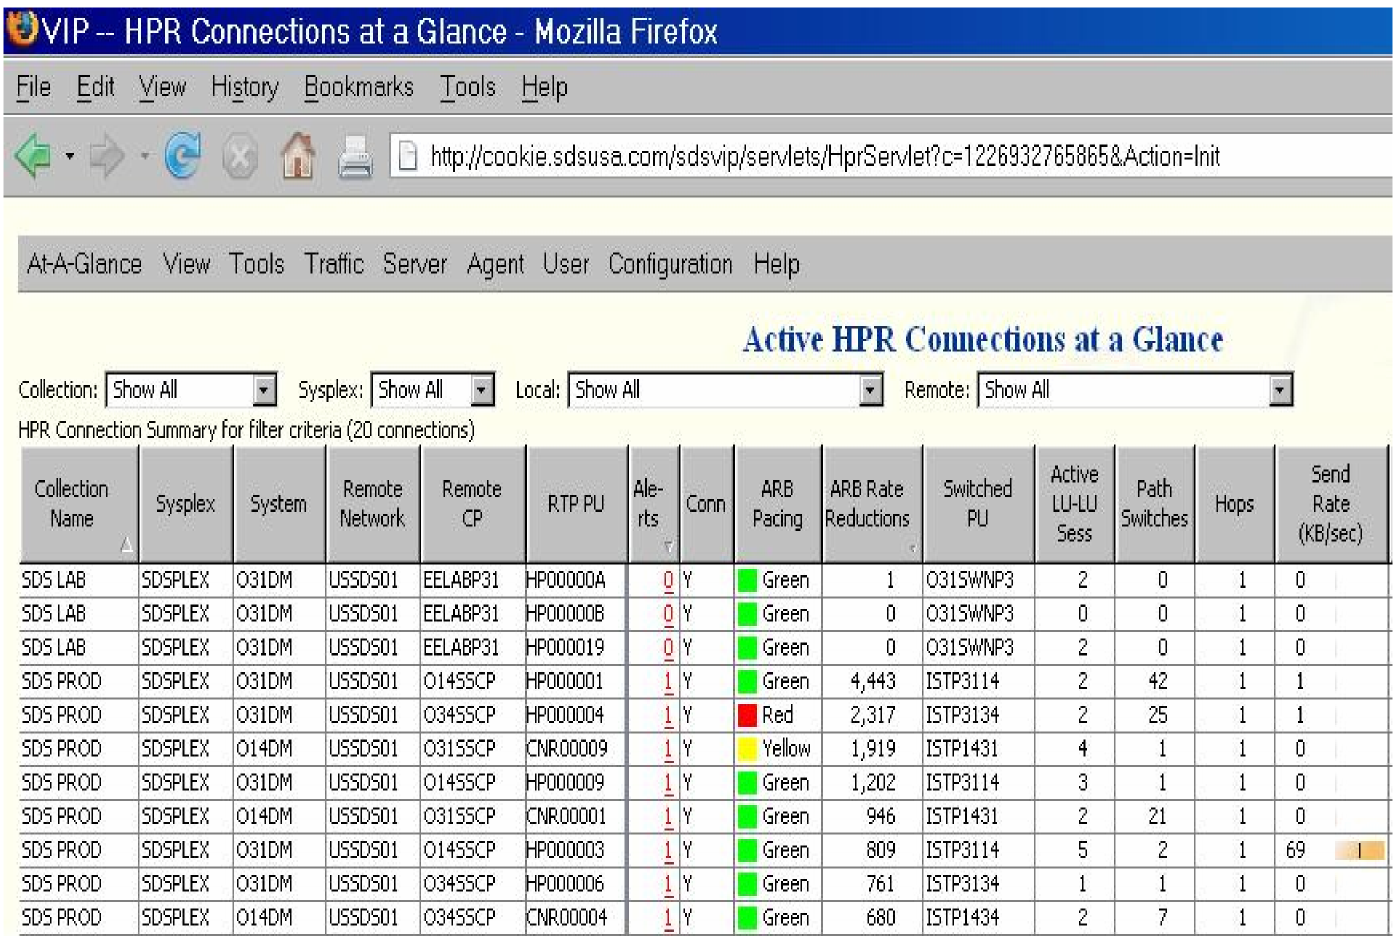

Collection Name	Sysplex	System	Remote Network	Remote CP	RTP PU	Ale-rts	Conn	ARB Pacing	ARB Rate Reductions	Switched PU	Active LU-LU Sess	Path Switches	Hops	Send Rate (KB/sec)
SDS LAB	SDSPLEX	O31DM	USSDS01	EELABP31	HP00000A	0	Y	Green	1	O31SWNP3	2	0	1	0
SDS LAB	SDSPLEX	O31DM	USSDS01	EELABP31	HP00000B	0	Y	Green	0	O31SWNP3	0	0	1	0
SDS LAB	SDSPLEX	O31DM	USSDS01	EELABP31	HP000019	0	Y	Green	0	O31SWNP3	2	0	1	0
SDS PROD	SDSPLEX	O31DM	USSDS01	O14SSCP	HP000001	1	Y	Green	4,443	ISTP3114	2	42	1	1
SDS PROD	SDSPLEX	O31DM	USSDS01	O34SSCP	HP000004	1	Y	Red	2,317	ISTP3134	2	25	1	1
SDS PROD	SDSPLEX	O14DM	USSDS01	O31SSCP	CNR00009	1	Y	Yellow	1,919	ISTP1431	4	1	1	0
SDS PROD	SDSPLEX	O31DM	USSDS01	O14SSCP	HP000009	1	Y	Green	1,202	ISTP3114	3	1	1	0
SDS PROD	SDSPLEX	O14DM	USSDS01	O31SSCP	CNR00001	1	Y	Green	946	ISTP1431	2	21	1	0
SDS PROD	SDSPLEX	O31DM	USSDS01	O14SSCP	HP000003	1	Y	Green	809	ISTP3114	5	2	1	69
SDS PROD	SDSPLEX	O31DM	USSDS01	O34SSCP	HP000006	1	Y	Green	761	ISTP3134	1	1	1	0
SDS PROD	SDSPLEX	O14DM	USSDS01	O34SSCP	CNR00004	1	Y	Green	680	ISTP1434	2	7	1	0

Figure 3: Just a fraction of the 42 columns of HPR data, per active connection, available with VIP's 'HPR Connections at a Glance' screen with its distinctive color-coded ARB congestion indications and outstanding alerts (in red) prominently in view. A horizontal scroll bar provides access to the other data which includes send/receive counts.

- An invaluable **EE traceroute** feature [Fig. 4], unique to VIP, that can dynamically test end-to-end EE path availability, from VTAM to a remote EE node – without the need for any manual intervention – and report back the IP addresses of all intermediate nodes, the time taken to traverse each hop and any packet losses that were experienced.
- A easy-to-master and accommodating **'Alerts at a Glance'** facility [Fig. 5], that not only clearly describes the cause of each alert but also provides details as to where in the network hierarchy the alert occurred and an explanation as to how that alert may be resolved. It is in essence an EE alerts wizard that ensures that no EE alerts go unresolved due to lack of visibility or the inability to locate information as to how to fix that type of problem.
- The ability to quickly access a comprehensive, unstinting log of **EE history** tabulations that spans availability details, data transfer statistics, congestion control patterns, path switch audit and queue length summaries. This information can even be displayed in graphical, bar chart format to study historic trends [Fig. 6].

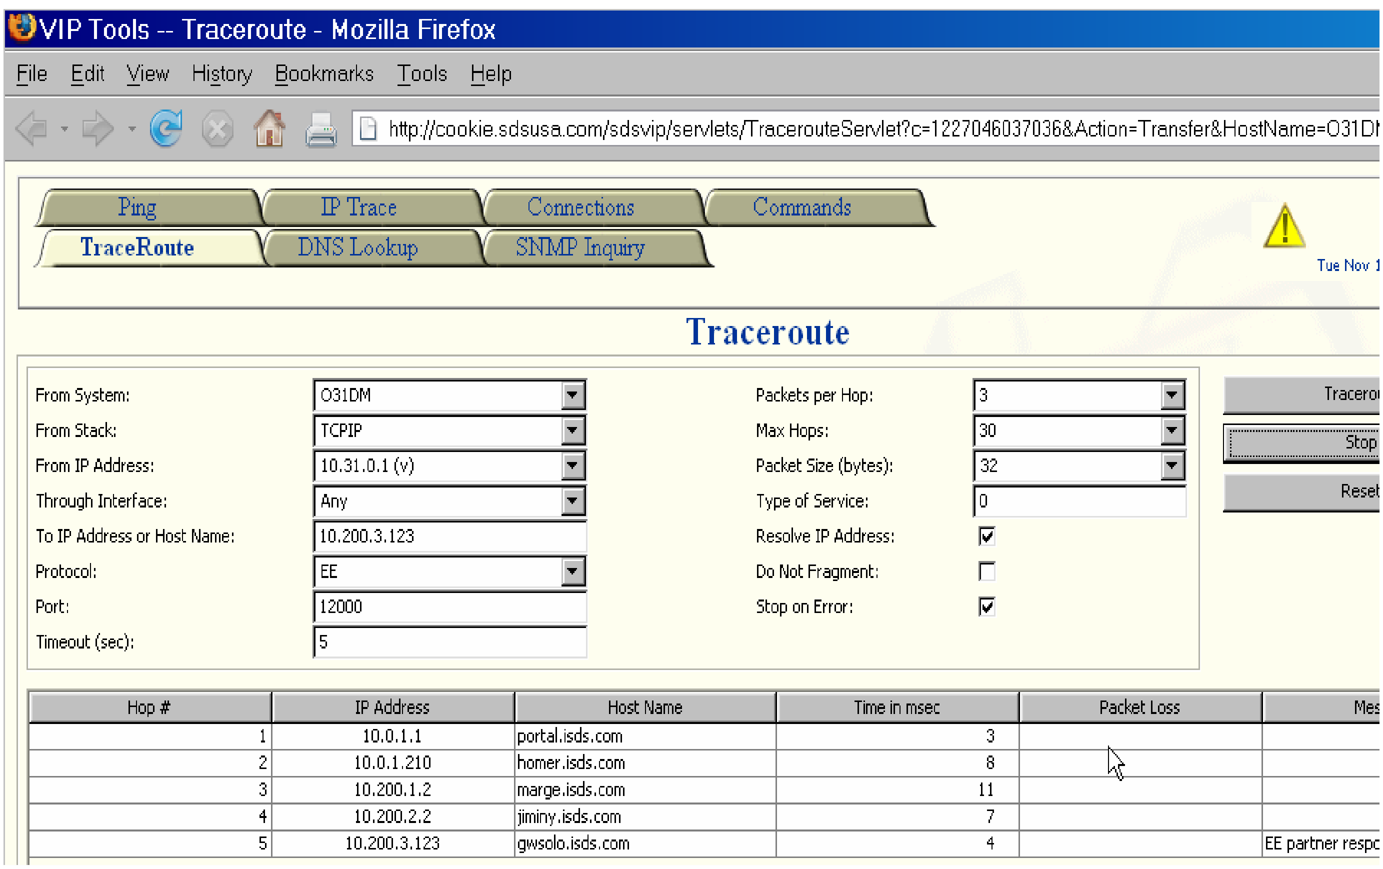

Figure 4: VIP's automatic IP address detecting, EE 'Traceroute' function that provides a detailed audit trail of an end-to-end EE connection in terms of the HPR hops involved, the time taken to traverse each hop and details of any packet losses incurred. [This screen was cropped, at right, for space reasons.]

Function Overlaps Necessitate Specialized Monitoring

HPR was not designed to run over UDP/IP, as is the case in EE. Instead, it was envisaged that HPR would work with Asynchronous Transfer Mode (ATM) – the 53-byte cell switching technology that IBM and the rest of the IT world were convinced, in the early 1990s, would be the ultimate in network transport mechanisms. ATM was a Layer 2 technology. Consequently, HPR's RTP and ARB components were designed to provide robust, Layer 3-4 functionality, comparable to that in TCP/IP.

Having HPR on top of UDP/IP à la EE results in the duplication of certain key functions – in particular dynamic rerouting (due to link level failures or congestion) and congestion control. HPR's rerouting, though automated and dynamic, cannot be performed by intermediate nodes. Only HPR end nodes can establish new paths and perform path switches. In today's IP networks dynamic rerouting and traffic throttling to control congestion is routinely done, transparently and efficiently, by routers and switches. In many instances IP routers will circumvent link-level issues well before any problems become apparent to HPR. [There is also a VTAM 'wait-prior-to-HPR-path-switch' timer to minimize potential 'race conditions' where HPR and IP could both be trying to fix the same problem.]

HPR does have the option of performing path switches to find a better (e.g., faster or less congested) route to an end destination. But, in the main, HPR path switches occur because an underlying problem has been detected in an existing path. So, when HPR is forced to perform a path switch, it typically indicates that something unusual has taken place, or for that matter, is continuing to take place in the underlying network – i.e., the IP network.

Ditto for HPR ARB congestion control. In most instances, if the underlying IP network is functioning as hoped, there should not be any undue HPR-level congestion. Thus, if there is significant ARB pacing, as denoted by VIP with a yellow or red semaphore [Fig. 3], you have no choice but to investigate – and do so quickly. As the saying goes, don't walk; run. Unnecessary ARB pacing by VTAM, due to what may be an easily fixed IP issue, could be squandering costly mainframe cycles.

As discussed above, VIP's 'HPR Connections at a Glance' [Fig. 3] provides detailed statistics on HPR path switches, congestion status and outstanding alerts. Moreover, there is VIP's EE traceroute [Fig. 4] that can then be immediately invoked, per connection, to determine the exact, *current* IP configuration of any given EE path. Thus, with VIP, IT professionals get unparalleled visibility into exactly what is taking place vis-à-vis HPR protocols and UDP/IP.

At this juncture it is also important to note that VIP is not just a management tool for HPR and EE. [See table on page 17] Instead, VIP is a highly proven, full-function, **real-time mainframe TCP/IP monitor** that is noted for its penetrating access into the inner workings of IP networks – for example, IP fragmentation interpretation, OSA(-E) usage statistics, IP stack analysis, SNMP MIB queries and network activity status.

Going back to the start of this White Paper and the reference to the swan -- with VIP, when it comes to EE, you get to see what is happening both above and below the surface. That in a nutshell is the 'take away message' of this White Paper. When it comes to staying on top of EE, VIP provides monitoring, management and optimization capabilities above and below the crucial UDP/IP boundary.

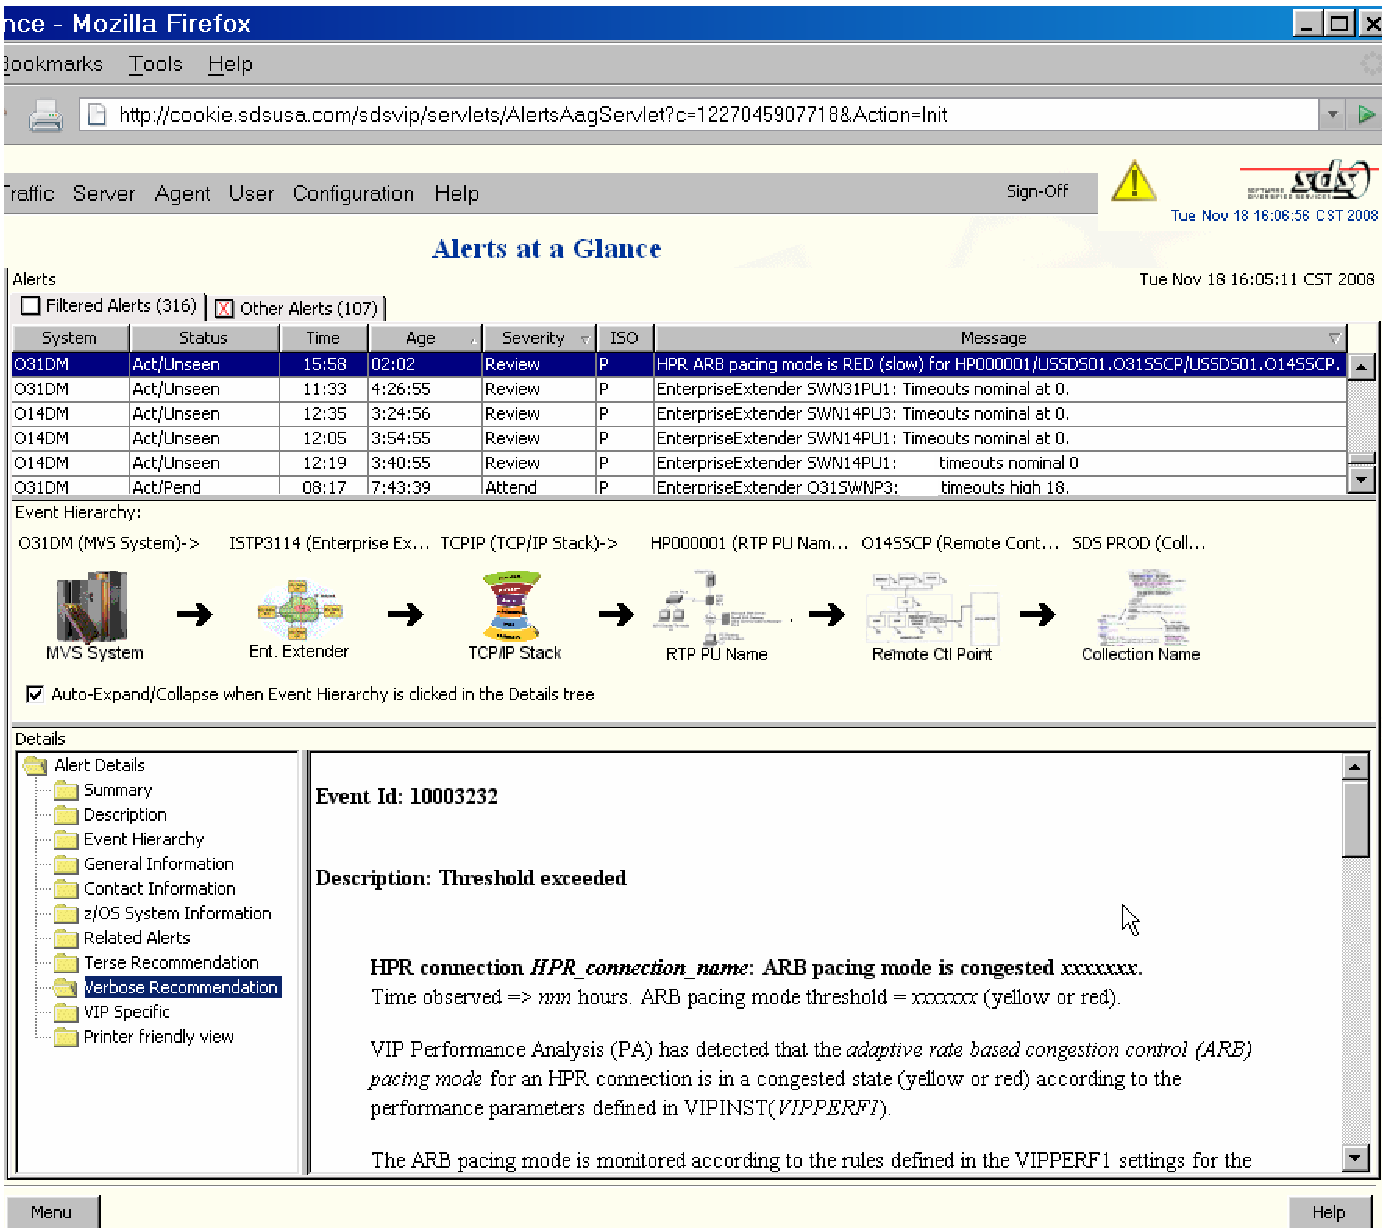

Figure 5: VIP's helpful, information-packed 'Alerts at a Glance' vis-à-vis EE. [cropped at left for space].

RESOLVING REAL-LIFE EE PROBLEMS WITH VIP

Now that we know, albeit but in overview, the EE-specific management functions in VIP, it is worthwhile to briefly examine how these capabilities are gainfully used, in real-life, by IT professionals. Let's start with a scenario where you get a call, from an EE end-user at a remote site, complaining of recent ***erratic performance*** – sometimes verging towards being unacceptable. The first thing you would obviously do would be to check the entry for that remote node/connection on the 'Enterprise Extender at a Glance' display [Fig. 3].

This 'At a Glance' display will immediately tell us:

- Pending ***alerts*** against that connection, which if non-zero, would need to be checked and resolved – typically via the 'Alerts at a Glance' display with its built-in fault resolution suggestions [Fig. 5].
- Observed ***availability*** for this connection – where a value less than 100% would denote that there has indeed been some intermittent outages, which could account for the erratic performance.
- Current ***hop count*** for this connection, indicating the number of IP intermediate nodes presently being traversed by this connection. If this number appears to be higher than expected based on your knowledge of the physical network you would definitely want to do some further investigation at this point. A larger than anticipated hop count could be the result of one or more HPR-level path switches. What we have to now establish is whether there has indeed been HPR-level path switches, and if so, the cause for these switches. Multiple HPR-level path switches, in the recent past, could definitely account for the reported erratic performance in two separate ways:
    1. An intermittent or permanent problem in the IP network, possibly with a router or OSA interface, that was disruptive enough to force HPR to perform a path switch. The erratic performance may have been due to packet loss or packet delay due to this IP layer problem.
    2. Noticeable pauses in end-to-end HPR data transmission when HPR is in the process of performing a path switch. The reported erratic performance could have been due to these HPR-level pauses, or a combination of these and the delays due to the IP network problems.

For our peace of mind, we also want to make sure while we are at it that any HPR path switches that may have occurred are not due to an 'inappropriate' setting on VTAM, given that HPR does have the option of performing path switches, routinely, in an effort to locate better paths – even when there are no problems occurring on the current path. So we need to check all of this out and try and quickly eliminate as many factors as possible. We may also want to get the problem end node on a more optimum, i.e., lower hop count, path because the longer path may be impacting response times, exacerbating the erratic performance.

The good news is that we have multiple VIP tools, right at our fingertips, to help us quickly drill-down to an additional vista of detail. We could, obviously, just do a right-click and invoke VIP's laser-like **EE traceroute** feature [Fig. 4] and get the lowdown on the current end-to-end path – including the current per hop delay and packet loss statistics. But we know the connection is active. So we should hold back on this for a few minutes and focus on using the available HPR path switch data to see if this will explain the erratic performance

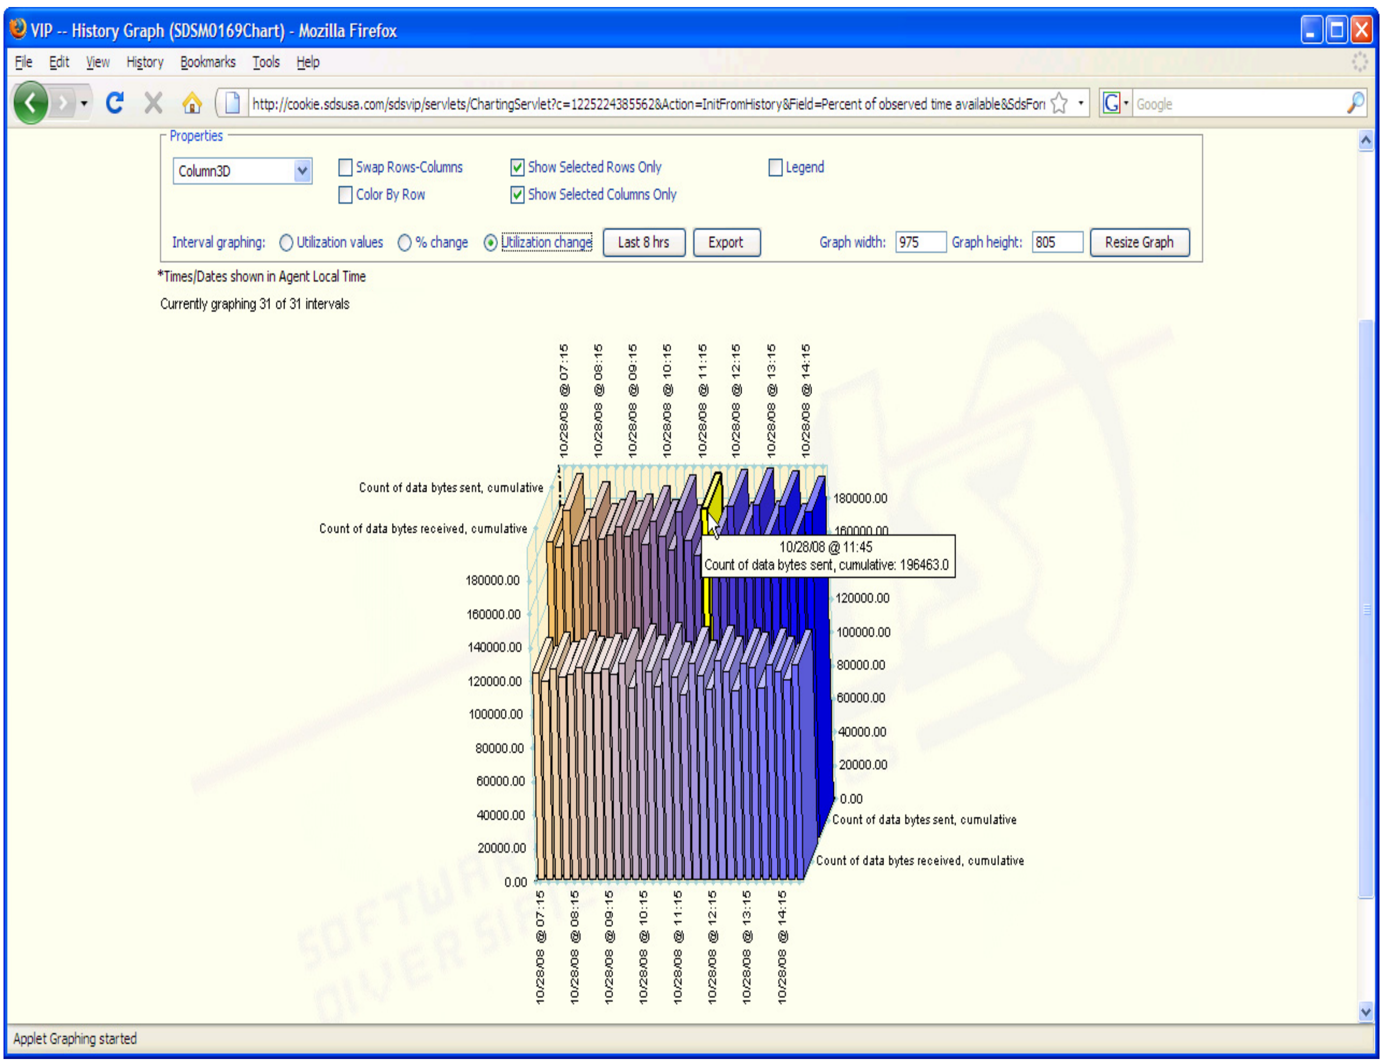

Figure 6: VIP's historical Enterprise Extender data display capability which can be customized to report data spanning days, weeks or months and can present this data in a number of compelling graphical formats with the histogram display shown here being but one option.

We should now bring up the 42-column ***'Active HPR Connections at a Glance'*** display [Fig. 3]. This display provides eight path switch specific fields of information [Fig. 7 below]– including the time of the last switch and the reason for that path switch.

Active HPR Connections at a Glance

Remote: Show All | Byte, NLP Scale: Auto

Fri Nov 14 12:31:31 CST 2008

ueue cvd egs	Queue Ack Wait	Ack Wait Max NLPs	Ack Wait Max Date	Ack Wait Max Time	P/S Local Init	P/S Remote Init	P/S Cause Errors	P/S Cause Retries	P/S Last Rsn	P/S Last Date	P/S Last Time	Local Network	Local CP
0	0	4	03Nov2008	10:35:53	0	0	0	0				USSDS01	O31SSCP
0	0	2	03Nov2008	10:35:58	0	0	0	0				USSDS01	O31SSCP
0	0	10	03Nov2008	10:35:58	0	0	0	0				USSDS01	O31SSCP
0	0	29	06Nov2008	09:18:39	5	33	5	0	S	14Nov2008	12:03:19	USSDS01	O31SSCP
0	3	19	06Nov2008	09:08:16	0	25	0	0	P	07Nov2008	08:24:12	USSDS01	O31SSCP
0	1	165	13Nov2008	10:11:41	1	0	1	0	S	07Nov2008	08:24:01	USSDS01	O14SSCP
0	0	137	13Nov2008	10:12:26	0	1	0	0	P	07Nov2008	08:24:13	USSDS01	O31SSCP
0	0	200	12Nov2008	09:27:14	2	0	2	0	S	06Nov2008	08:48:00	USSDS01	O31SSCP
0	0	14	06Nov2008	08:42:55	12	5	12	0	P	P: Partner initiated last path switch	[illegible]	SDS01	O14SSCP
0	0	16	06Nov2008	09:18:29	0	5	0	0	P	13Nov2008	13:33:27	USSDS01	O14SSCP
0	0	6	13Nov2008	19:31:18	0	2	0	0	P	06Nov2008	08:48:02	USSDS01	O14SSCP
0	0	3	13Nov2008	12:37:37	0	0	0	0				USSDS01	O14SSCP
0	0	4	01Nov2008	12:09:49	0	0	0	0				USSDS01	O14SSCP
0	0	3	01Nov2008	12:09:49	0	1	0	0	P	05Nov2008	15:06:44	USSDS01	O31SSCP
0	0	1	05Nov2008	15:06:44	0	0	0	0				USSDS01	O31SSCP
0	0	2	01Nov2008	11:56:14	0	0	0	0				USSDS01	O31SSCP
0	0	1	13Nov2008	13:31:18	0	0	0	0				USSDS01	O14SSCP
0	0	2	01Nov2008	11:56:14	0	0	0	0				USSDS01	O14SSCP
0	0	6	13Nov2008	09:53:59	0	0	0	0				USSDS01	O14SSCP
0	0	4	13Nov2008	09:53:59	0	0	0	0				USSDS01	O14SSCP
0	0	1	13Nov2008	09:55:15	0	0	0	0				USSDS01	O14SSCP
0	4	629			20	72	20	0					

< ------------ "P/S" – PATH SWITCH RELATED FIELDS ------------ >

Figure 7: Seven of the 'Path Switch' specific fields from VIP's 'Active HPR Connections at a Glance' display with the 'number of path switches per connection' field appearing earlier [i.e., to the left] in the display as shown in Fig. 3. The "P/S Last Rsn" field indicates the reason for the last path switch.

The 'Path Switches' field [Fig. 3 & 8] will tell us the number of switches that have taken place for this end-to-end connection since it was established. If this number appears to be large, especially relative to those for other connections, we could definitely be on the right track here. The 'P/S Last Date' and 'P/S Last Time' [Fig. 7] will now tell us when the last path switch occurred. If the last switch occurred shortly prior to the end-user's call, it will confirm that we are making sound progress. The 'P/S Last Rsn' code will give us our next clue. An 'S,' indicating 'short request retry limit exhausted,' would indicate an intermittent IP level problem (or possibly a VTAM timer that needs to be 'upped'). An 'A,' indicating that the last switch was a unilateral, automatic attempt on the part of VTAM to search for a better route might mean we need to check the VTAM **PSRETRY** option to determine whether we should throttle back on these automatic searches.

The 'HPR at a Glance' also tells us, unmistakably, the ***congestion status*** for that connection – with its green, yellow and red semaphores [Fig. 8 below]. A red or yellow designation for this link would also explain the erratic performance – as could any unresolved alerts. Congestion problems would also be reflected by a high 'ARB Rate Reduction' count – as with the fourth entry in Fig. 8 below. That connection also has a high 'Path Switches' count. If this is the connection we are investigating, then it is not surprising that we got a call about erratic performance. But we can also see, thanks to VIP, that the connection is performing quite well right now – there are no queued packets and the send rate is relatively robust.

Active HPR Connections at a Glance

: Show All Remote: Show All Byte, M

Fri N

:TP PU	Ale-rts	Conn	ARB Pacing	ARB Rate Reductions	Switched PU	Active LU-LU Sess	Path Switches	Hops	Send Rate (KB/sec)	Smthd RTT (ms)	NLPs Queued	NLPs 5 Min Rcvd	NLPs 5 Min Sent
0000A	0	Y	Green	1	O31SWNP3	2	0	1	0	247	0	0	0
0000B	0	Y	Green	0	O31SWNP3	0	0	1	0	165	0	0	0
0000C	0	Y	Green	0	O31SWNP3	2	0	1	0	22	0	0	0
00001	1	Y	Red	3,686	ISTP3114	2	37	1	3	26	0	354	372
00004	1	Y	Green	1,987	ISTP3134	2	25	1	0	64	0	191	190
:00009	1	Y	Red	967	ISTP1431	9	1	1	0	18	0	283	274
00009	1	Y	Yellow	921	ISTP3114	9	1	1	0	26	0	274	282
00003	1	Y	Green	804	ISTP3114	6	2	1	0	23	0	6	11
:00001	1	Y	Red	546	ISTP1431	2	16	1	3	27	0	371	354
:00004	1	Y	Green	368	ISTP1434	2	5	1	1	56	0	201	191
:00003	1	Y	Green	16	ISTP1431	6	2	1	0	38	0	11	6
:0002E	1	Y	Green	2	ISTP1434	2	0	1	0	49	0	1	1
:00006	0	Y	Green	11	ISTP1434	6	0	1	0	38	0	8	6
00006	0	Y	Green	1	ISTP3134	1	1	1	0	36	0	2	2

Figure 8: An expanded view of the 'Alerts,' 'ARB Pacing,' 'Send Rate' and 'NLP' (Network-Layer Packets) fields of the 42-column, 'Active HPR Connections at a Glance' display – complementing the other views shown in Figures 3 & 7.

By now, *though we have only spent less than a couple of minutes on this problem*, we already have some major clues and a plethora of actual data as to what has been taking place on this connection. At this juncture we may decide to invoke VIP's **EE traceroute** feature [Fig. 4]. This will give us an accurate picture of the configuration and status of the current end-to-end path replete with delay and packet loss statistics for each hop. We could, also using this same feature, perform an IP Trace or an SNMP MIB Query on an IP router or an OSA(-E) interface.

If ARB-level congestion is an issue we may also want to check for any 'rogue' IP fragmentation that may be a contributing factor. VIP's potent and intelligent IP **fragmentation monitor** [Fig. 9] will provide us with an in-depth review of any and all fragmentation at the UDP/IP level and even activate the necessary traffic analysis to decisively locate the source of the fragmentation.

But before we go any further and start exploiting VIP's arsenal of IP-related management tools to resolve this issue, what we have to ask ourselves is: *'How would we have found any of this EE and HPR specific data, such as hop counts, path switches, ARB congestion, EE alerts etc., without VIP?'*

Without VIP there would have been no visibility as to what was happening above UDP/IP. From an EE perspective we would have been flying blind. Since EE is invariably used to sustain business-critical SNA applications, in many instances in inter-company SNI configurations, not having the tools to monitor/manage EE could prove to be shortsighted and difficult to explain to 'C-level' management who are agitated that there is a major network outage impacting revenue.

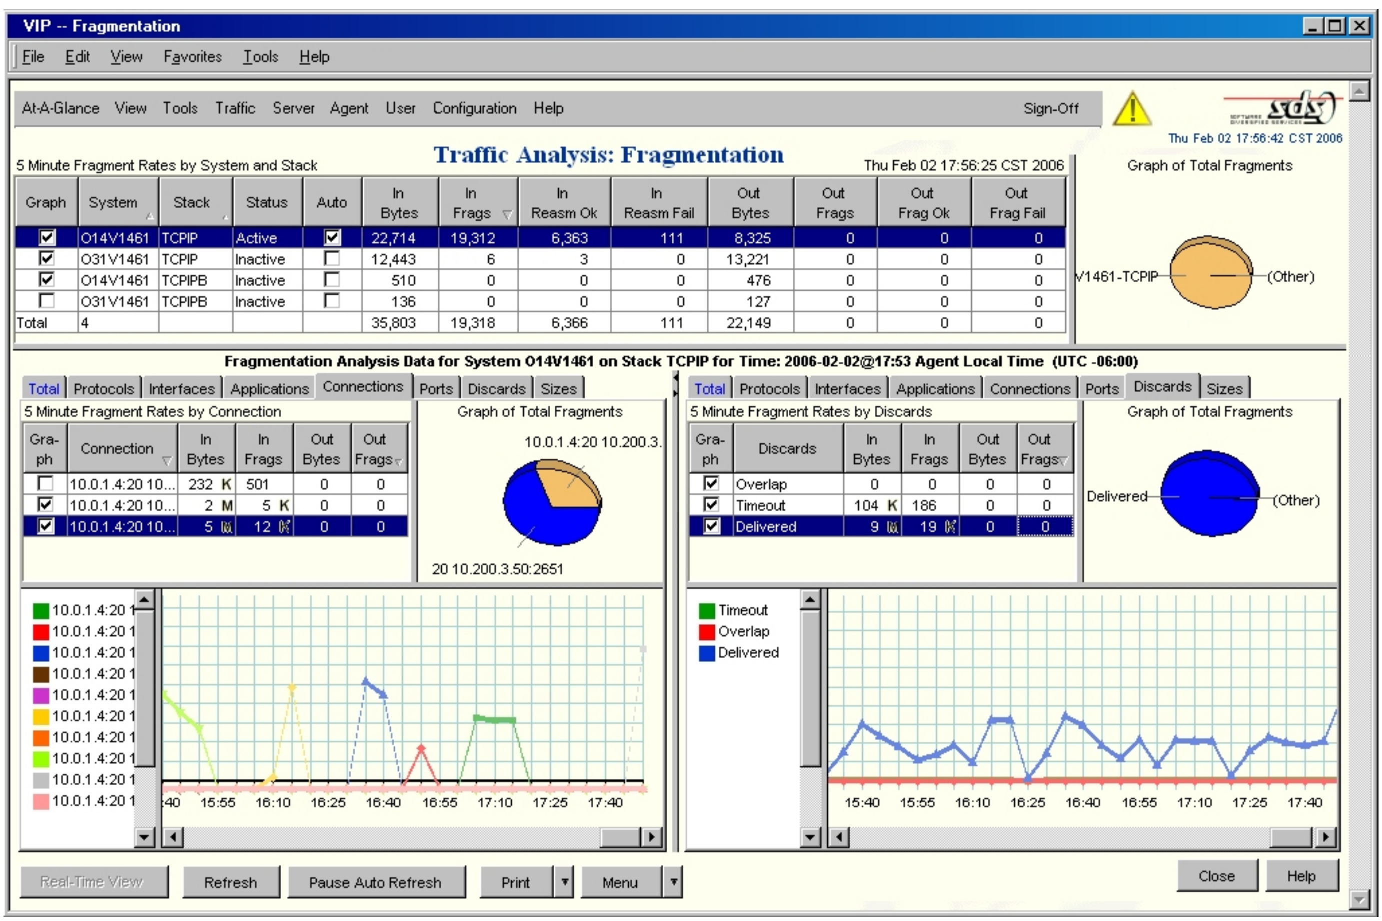

Figure 9: VIP's much admired and widely used IP fragmentation analysis feature.

The table on page 17 attempts to summarize the repertoire of EE/HPR, UDP/IP, TCP/IP, tn3270(E), HTTP, application, OSA, SNMP, FTP and Sysplex related monitoring and diagnostic capabilities available in VIP for efficacious mainframe-centric network management. The take-away message here being that with VIP there are no blind spots, let alone any black holes, vis-à-vis full spectrum, network management.

So now let's go back to the EE erratic performance problem we are trying to resolve. From what we learned from VIP's EE/HPR tools the underlying cause of the erratic performance was most likely an IP-level problem – but one that could have been exacerbated by HPR reacting to the problem with path switches and ARB congestion control. Whether IP fragmentation is the issue can be determined using VIP's fragmentation analysis. If it is, VIP will help us locate the source of the fragmentation so that we can fix that. If fragmentation isn't the issue we could invoke VIP's 'network-radar-like' **'Network at a Glance'** screen [Fig. 10]. This screen will tell us, at once, about any outstanding issues across the network. We will want to pay particular attention to the status and alerts pertaining to IP stacks, interfaces, remote hosts and OSA.

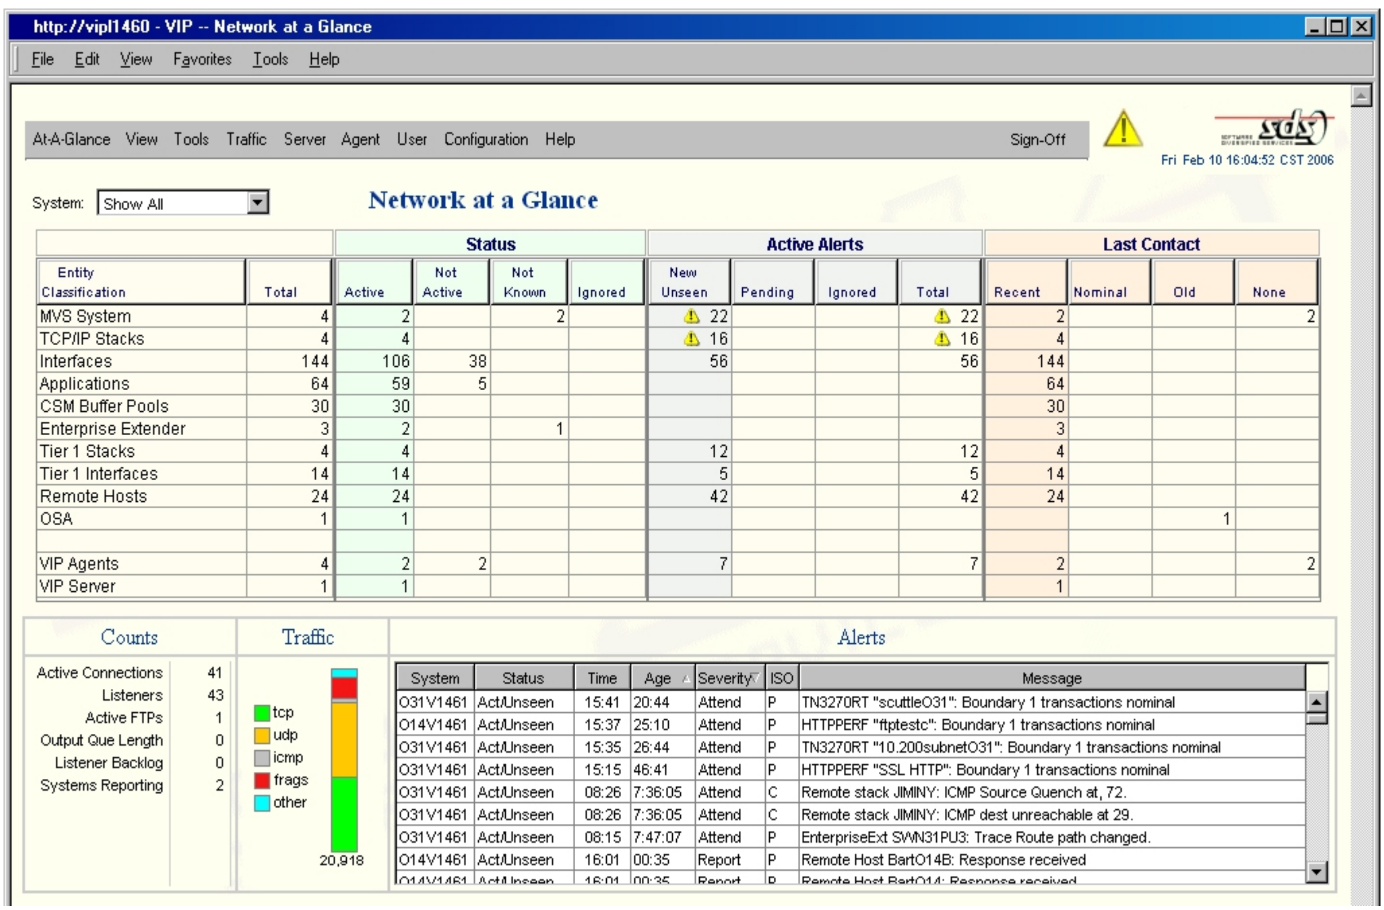

Figure 10: VIP's 'Network at a Glance' screen that provides real-time data on all network related resources.

To describe how the plethora of monitoring and diagnostic tools available within VIP can be used to isolate and resolve IP-level network problems is beyond the scope of this EE-oriented White Paper. Online, interactive demos, freely provided by SDS, are indubitably the best way to get to know these capabilities. With these online demos you can go through an actual sequence of actions, such as checking the MIB of a remote router that may be the culprit that we are trying to locate, and see how easy it is with VIP to sequentially drill down, deeper and deeper.

The bottom line here is that monitoring and managing EE has to be a two-prong approach with both HPR-specific and IP-related diagnostics. Having one without the other is not going to work. Hence, the irrefutable advantage of VIP's integrated, full spectrum solution for EE.

VIP's Agent/Server 'Green' Architecture

VIP uses an extremely efficient, 'lean-and-mean' agent/server architecture [Fig. 11], with VIP monitoring agents available for all the relevant releases of z/OS. Data gathered by the agents are fed, in real-time, to one or more VIP servers [with the agents making sure that only the data that has changed is sent to the servers so as to minimize agent-server traffic]. Deploying multiple servers guarantees resilience for 'zero downtime' operations. The VIP servers, which are implemented in Java, can be deployed on PCs running Windows, on Linux/Unix servers [including Linux LPARs], or on an 'MVS' LPAR with Unix System Services (USS).

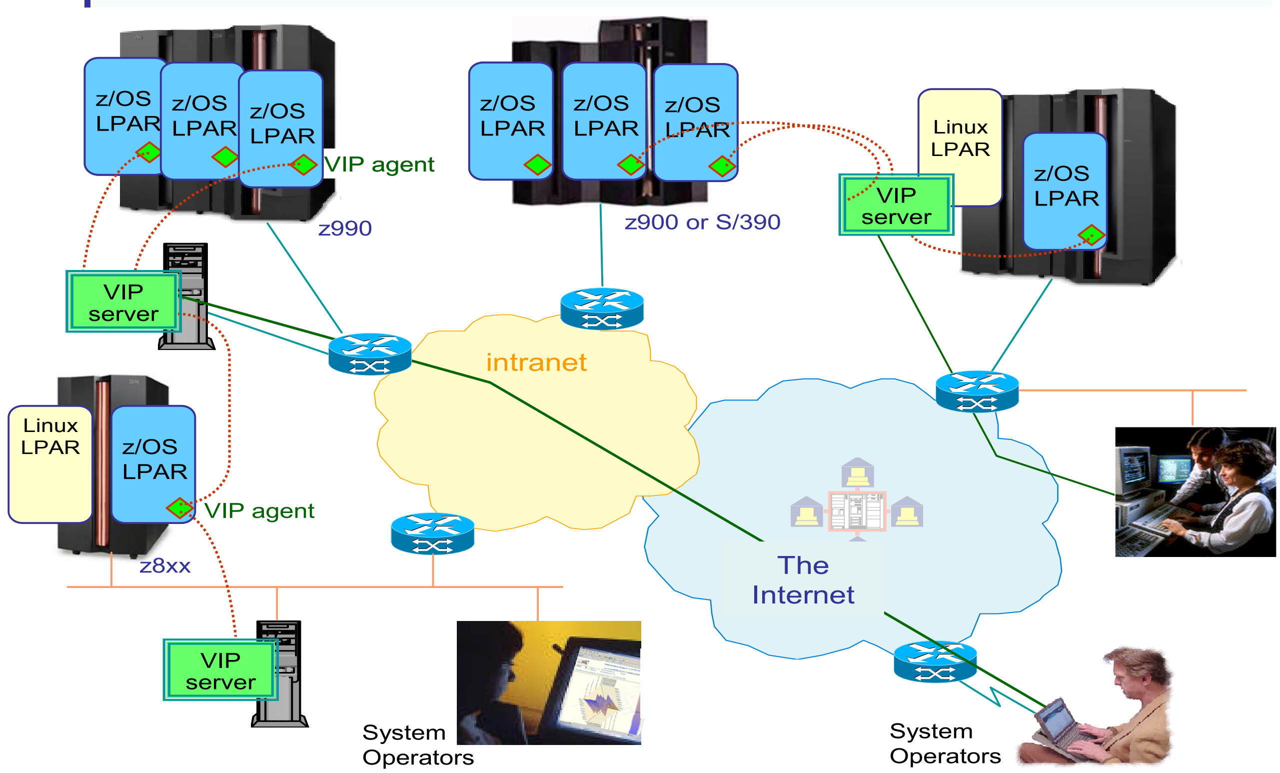

Figure 11: VIP's flexible, 'light-on-resources,' mainframe off-load agent/server architecture

A *Partial* Repertoire of VIP's Network Monitoring and Management Capabilities

EE and HPR Specific

- EE connections & ports at a Glance [Fig. 2].
- HPR connections at a Glance [Figs. 3, 7 & 8].
- EE trace-route [Fig. 4].
- EE alerts at a Glance [Fig. 5].
- EE History [Fig. 6].

IP Related

- Network at a Glance [Fig. 10].
- Activity at a Glance.
- IP packet trace.
- IP stack configuration and utilization (tier-1 and remote).
- IP traceroute (TCP or UDP).
- Alerts at a glance
- IP fragmentation analysis [Fig. 9].
- Ping (ICMP & UDP).
- SNMP MIB queries.
- Remote hosts at a Glance.
- History charting.

TN3270(E) & TELNET

- With SNA LU name correlation.
- RFC 2562 compliant response time monitoring (RTM).
- Session history.

TCP & HTTP Applications

- Applications at a Glance with CPU/buffer utilization, round-trip times, etc.
- HTTP performance and response time monitoring.

FTP

- FTP at a Glance.
- FTP history.

OSA(-E)

- OSA at a Glance.
- Support for both SNA & IP traffic.

Additional Functionality

- Integrated DNS lookup.
- Sysplex Distributors at a Glance.
- Intrusion Detection data.
- IBM SMF records.
- Connections Explorer.
- Fully exploits IBM's Network Management Interface (NMI).
- Monitoring of system console messages and control blocks.
- Submit MVS system commands through VIP.

The VIP monitoring agents are written in optimized IBM assembler for maximum speed and efficiency. It is these agents that access the various mainframe resources for management data and provide the data for diagnostics. Agents access EE, HPR, VTAM, IP stack and z/OS data via IBM's strategic **Network Management API (NMI)**. Thanks to VIP's agent/server architecture, LPARs or mainframes can be easily and quickly added, non-disruptively, to an existing VIP-managed network.

VIP's easy-to-master, highly graphical user interface is Web browser-based. Thus, *authorized* operators can readily access all VIP services from any network-attached client running a standard Web browser as depicted in Fig. 11. VIP's operator authorization is based on IBM's z/OS System Authorization Facility (SAF). Via SAF, the VIP features available to any given operator can be selectively enabled or disabled.

Meaningful, easy-to-comprehend data presentation is another trademark of VIP. Careful thought has been given to the screen designs. Thus VIP screens don't overwhelm operators with 'unstructured' data. Instead graphs, tables, semaphores and icons are extensively used to present a clear and concise picture of network status and operation. Operators, however, always have the option to quickly 'drill-down' to the next layer of detail – as was discussed in the above section.

VIP can readily interoperate with central automation products such as HP's OpenView and IBM's NetView®. VIP alerts can be automatically forwarded to such central automation products, sent to a specified system console or used to generate an SNMP trap. It is also important to point out, at least in passing, that:

- VIP is easy to install and maintain.
- VIP is very reliable and meets the 'non-stop' criteria of mainframe operations.
- VIP is well integrated with z/OS and does not impact or interfere with any of the mainframe software components.
- VIP uses standard IBM z/OS APIs, protocols and guidelines and as such is very z/OS conformant.
- SDS' technical support for VIP is very professional, competent and responsive.
- VIP offloads network monitoring related data processing, data analysis and data presentation so that these functions do not consume resources that could instead be used for mainframe production workloads.

The Bottom Line

EE, 'HPR-over-UDP/IP,' now represents SNA's final frontier – the IBM endorsed, strategic means for sustaining business critical SNA/APPN applications across IP networks, without sacrificing SNA's native COS and routing. EE, however, unlike its illustrious forebears, SNA and APPN, was not a painstakingly designed, end-to-end architecture. Instead it was mashup – a decade before mashups became fashionable. In an effort to enable SNA across IP, IBM, working with many of the router vendors, came up with EE – layering HPR on top of UDP/IP. The result, as shown in Fig. 1, was a rather elongated protocol stack with considerable overlap of functionality between HPR and IP.

Given that HPR, independent of IP, has the option to perform end-to-end path switches and congestion control, one cannot, by any stretch of the imagination, adequately manage a business-critical EE network just using TCP/IP oriented network management products. You have to have specialized HPR support as is available with VIP with its well defined suite of EE/HPR capabilities.

The level of HPR data made available by VIP, spanning hop counts, path switches, ARB congestion, transmission rates and NLP queuing statistics, is unprecedented. This degree of data, at that granularity, can only be obtained in real-time via the proficient exploitation of IBM's NMI. Any other technique, such as packet tracing, will prove to be woefully inadequate. VIP's EE traceroute and alert reconciliation, furthermore, provide incisive EE/HPR-specific means to quickly isolate and resolve EE related issues.

But the real beauty here is that VIP, despite these industry-leading EE/HPR capabilities, is not just a 'point solution.' Instead, it is a highly proven, full-spectrum mainframe TCP/IP monitor with a highly effective 'green' architecture. It includes comprehensive, '*no-blind- spots-whatsoever,*' support for all the components that make up the network – encompassing everything from applications to the remote routers.

Just keep in mind the image of the swan. When it comes to EE monitoring/management you have to have visibility and access to what is happening both above and below the 'surface' – where the 'surface' in this case is the UDP/IP boundary. VIP, as you now realize, really does work both above and below this crucial boundary.

With VIP it is thus guaranteed that you can stay on top of EE – at all times.

Selected Glossary

ANR	Agile, low-overhead, Layer 2, connectionless routing mechanism used by HPR which is based upon using a Routing Information Field (RIF) à la Token Ring LAN source-route bridging.
APPN	A landmark 1986 rework of the original SNA architecture to enable plug-and-play, peer-to-peer networking unhampered by centralized control from a mainframe.
ARB	HPR's data flow control mechanism for proactive congestion control which is a sub-component of HPR's RTP.
Control Point	SNA/APPN/HPR functionality that performs authorization, directory services and configuration management.
COS	SNA's Class of Service traffic prioritization scheme.
DLSw	Widely used SNA/APPN(/NetBIOS)-over-TCP/IP transport mechanism which, however, unlike EE does not support SNA COS or routing.
DLUR	An IBM mechanism to support traditional SNA traffic (designed to be dependent on mainframe control) across peer-to-peer APPN or HPR networks.
EE	HPR-over-UDP/IP, created by committee and codified in RFC 2353 in 1998, which permits SNA/APPN networking, replete with native COS and routing, across IP networks.
HPR	An SNA architecture developed by IBM in the early 1990s to imbue SNA/APPN with dynamic alternate routing, nimble intermediate node routing and proactive congestion control in order make SNA networking more competitive with TCP/IP.
IP	The primary, underlying, connectionless protocol which is the basis for all Internet Protocol Suite based networking.
LLC-2	Connection-oriented version of the Layer 2 Logical Link Control protocol.
LU	SNA's software interface (or 'port') through which end users gain access to the SNA network.
LU 6.2	SNA's protocol suite for program-to-program communications.
NMI	IBM's Network Management Interface for obtaining key management related data in real-time.
Ping	Widely used protocol for checking the availability of a remote host on an IP network.
RTM	Response time monitoring scheme for 3270s which is now supported by tn3270(E) per RFC 2562.
RTP	End-to-end, full-duplex transport mechanism used by HPR to realize reliable, in-sequence data delivery with selective retransmission

	capabilities to cope with error recovery. RTP functions, which also include packet segmenting and ARB-based congestion control, are not performed at HPR intermediate nodes. RTP, in essence, is EE's TCP layer.
SAF	Security Authentication Facility; e.g., RACF, ACF 2 or TopSecret.
SMF	IBM's long-standing System Management Facilities mechanism for collecting mainframe-related management data.
SNMP	Set of TCP/IP-centric, network management protocols.
tn3270(E)	Widely used, client-server technology that permits TCP/IP clients to access mainframe resident SNA applications using 3270 data stream.
UDP	Connectionless, transport layer protocol that is a part of the Internet Protocol Suite.
VTAM	IBM's mainframe-based software for implementing an SNA/APPN/HPR node within an LPAR.

Software Diversified Services

Software Diversified Services (SDS), [www.sdsusa.com] based in Minneapolis, MN, has been providing premium mainframe solutions to the IBM world since 1982. It currently has in excess of 1,000 mainframe customers worldwide.

SDS' mainframe product repertoire now includes over twenty z/OS, VM and VSE products, with the highly regarded Vital Signs VisionNet (VIP) IP Monitor being one of these. SDS also markets PC software related to mainframe operations. The products marketed by SDS focus on network management, performance monitoring, report distribution and data compression.

SDS is noted for having the highest quality software, documentation, and technical support in this industry sector. SDS technical support has been rated #1 by the prestigious IBEX Bulletin.

The Author

Anura Gurugé [www.guruge.com] is an ex-IBMer who used to be 'Mr. SNA' for the longest time. His first book, *'SNA: Theory and Practice'* was published in 1984. He was heavily involved with Token-Ring switching, APPN, Frame Relay and Web-to-host.

These days he is a consultant and a professional writer. His latest book, critically acclaimed, is *'Popes and the Tale of Their Names.'*

He can be contacted at (603) 455-0901 or anu@wownh.com.