


Acuitive, Inc.



Virtual Resource Management

Industry Update and Analysis

Part of the Acuitive VRM Research Report subscription series.

March 2000

TABLE OF CONTENTS

1	<u>BRIEF INTRODUCTION</u>	4
1.1	<u>What Is Virtual Resource Management?</u>	4
1.2	<u>About This Report</u>	4
1.3	<u>Why I Care About VRM</u>	4
2	<u>1999 IN REVIEW</u>	6
2.1	<u>The Big Picture (It's Pretty Exciting)</u>	6
2.2	<u>Now You Can Start To Measure Success</u>	6
2.3	<u>Summary of Key Technical Trends</u>	9
2.3.1	<u>A New Taxonomy of Product Types</u>	9
2.3.2	<u>Cookies for Persistence</u>	9
2.3.3	<u>Bypass for Show, Delay for Dough</u>	12
2.3.4	<u>Getting Multi-Site Right?</u>	12
2.3.5	<u>Firewall Load Balancing – It's Not Just For Breakfast Any More</u>	13
2.3.6	<u>Caching and Server Load Balancing: Complementary or Competitive?</u>	14
2.3.7	<u>The Myths And Legends Of Security</u>	16
3	<u>Y2K MARKET DYNAMICS</u>	17
3.1	<u>Key Success Factors For VRM Vendors</u>	17
3.1.1	<u>Just "Be There"</u>	17
3.1.2	<u>SLB "Break Out" Strategy</u>	18
3.1.3	<u>Clear and Unique Positioning</u>	18
3.2	<u>Positioning Summaries</u>	19
3.3	<u>Hoover Projections</u>	20
3.4	<u>My Projections Are Wrong – Here's (Possibly) Why</u>	22
3.4.1	<u>More First Time Buyers (Or Less Vendor Loyalty)</u>	22
3.4.2	<u>F5 Networks Makes A Major Misstep</u>	22
3.4.3	<u>The Enterprise Market Takes Off</u>	22
3.4.4	<u>Cisco Gets Their Act Together</u>	22
3.4.5	<u>Colo Implementation Shift</u>	23
3.4.6	<u>Shift In Buyer Background</u>	23
3.4.7	<u>Market Consolidation</u>	24
4	<u>VENDOR UPDATES</u>	27
4.1	<u>Alteon</u>	27
4.1.1	<u>New News since May, 1999</u>	27
4.1.2	<u>Updated Analysis</u>	27
4.2	<u>ArrowPoint</u>	30
4.2.1	<u>New News Since May, 1999</u>	30
4.2.2	<u>Updated Analysis</u>	32
4.3	<u>Cisco</u>	34
4.3.1	<u>New News Since May, 1999</u>	34
4.3.2	<u>Updated Analysis</u>	35
4.4	<u>Coyote Point</u>	36
4.4.1	<u>New News Since May, 1999</u>	36
4.4.2	<u>Updated Analysis</u>	36
4.5	<u>F5 Networks</u>	37
4.5.1	<u>New News Since May '99</u>	37
4.5.2	<u>Updated Analysis</u>	40
4.5.3	<u>Prospects</u>	41
4.6	<u>Foundry</u>	42
4.6.1	<u>New News since May, 1999</u>	42
4.6.2	<u>Updated Analysis</u>	43
4.7	<u>CyberIQ Systems (The Company Formerly Known as Holontech)</u>	44
4.7.1	<u>New News Since May, 1999</u>	44
4.7.2	<u>Updated Analysis</u>	44
4.8	<u>HydraWeb</u>	46
4.8.1	<u>New News Since May, 1999</u>	46
4.8.2	<u>Updated Analysis</u>	46
4.9	<u>IBM</u>	47
4.9.1	<u>New News since May, 1999</u>	47

4.9.2	Updated Analysis	47
4.10	IPivot/Intel	48
4.10.1	New News Since May, 1999	48
4.10.2	Updated Analysis	48
4.11	RadWare	50
4.11.1	New News Since May, 1999	50
4.11.2	Updated Analysis	51
4.12	Resonate	53
4.12.1	New News since May, 1999	53
4.12.2	Updated Analysis	54
4.13	Others	57
4.13.1	The "F5 Enabled" Layer 4 Switches	57
4.13.2	Lightspeed Systems	57
1.	The gateway	57
4.13.3	Phobos	58
4.13.4	StoneSoft	58
4.13.5	Rainfinity	60
4.13.6	WebManage	61
5	SOME BACKGROUND INFORMATION	62
5.1	Who Is Acuitive?	62
5.2	About Our Reports and Subscription Service	62
5.3	About The Author	63
5.4	Thanks From The Author	63

1 Brief Introduction

1.1 What Is Virtual Resource Management?

Acuitive uses the term Virtual Resource Management (VRM) for the methods and procedures associated with making multiple networked devices “appear” to users and related components as one larger device. Want to make six NT servers and two Solaris servers appear as one large web server? VRM. Want to make three routers look like one large one? The answer is VRM. Want to make multiple firewalls, VPN devices, or proxy cache servers look like one big device? You should consider VRM.

1.2 About This Report

This report is part of the Acuitive VRM subscription program, kicked off with the publication of the following VRM documents published in May, 1999:

- Virtual Resource Management (VRM): Key Technologies, Tricks of the Trade, and Application Requirements (Version 3.0)
- Virtual Resource Management (VRM): Which Vendor is Right For You? (Version 3.0)

However, this document does not replace the previous documents. The nature of the coverage of the subject is quite different in this document. Previously, my target readership was entirely the technical decision makers involved with using the technology. While that readership is still an important target, I have found in the last 7-8 months that other communities are thirsty for information about this industry as well – product planners, vendor strategists, industry analysts, financial analysts, press, venture capitalists, investment bankers, investors, business development managers at large companies assessing M&A activity, and my mom.

Therefore, this document provides only a summary update of each vendor’s product line. Instead, the focus is on the state of the VRM industry and market, the immediate future and prospects for this market, and the nature of each vendor’s business, at least as well as I can understand and articulate it. In this spirit, and in order to drive my personal philosophy of “what good are friends if you can’t kick them once in awhile,” I have happily provided unsubstantiated (and therefore unarguable) personal opinion. To cap my arrogance off completely, I have also offered advice to many of the major vendors, which if they take, will either lead them to heightened success or it won’t.

1.3 Why I Care About VRM

I get asked all the time why I spend so much time looking at the VRM market and technology, because at first blush it seems like such a small niche the overall scheme of things.

My answer is partially because it is fun and partially because it is important.

Technically speaking, it is fun for me because VRM represents a form of technology convergence. I always like it when worlds collide. Convergence is a term usually used with respect to integration of voice and data, but I think it is equally valid to use the term in situations where networking technologies and practices collide with application/systems technologies and practices. Suddenly, two communities of people who in the past could ignore (and make fun of) one another have to find a language to communicate. If that’s not fun to watch I don’t know what is.

Another reason it is fun is that I really like most of the people I know and work with at each of the vendor companies. In general, these are the same kind of aggressive shaker/mover types who created the router industry. It is fun to work with these people when they are being mature and objective, and even more fun to observe them and call them on it when they are not. Potentially, this could lead to some serious situations, since I am observing things that they take seriously and are related to their livelihood. But I have found that this community of people is willing to forgive and forget as long as I demonstrate a passion for what they are doing, objectivity in my analysis, and devotion to doing my homework. So – things have sailed pretty smoothly.

The importance of this technology is another factor in my decision to spend time on it. As a \$270M market in 1999, it is not huge. But when you take into account that in 1996 it was close to a \$0 market, it seems pretty big.

Independent of the market size, I believe that VRM is fast becoming a highly visible and influential technology in the technical Internet community, and will soon have a position in the enterprise space similar to what routing and management platforms have had in the past. Once you have your router/switch network in place to move bits and packets around, the next thing you need to worry about is the relationship of end users to your applications. That is where VRM comes in. It overlays the network to provide a higher level of service level control. Networks are about connectivity. VRM is about the reliability, scalability, and smooth operation of the applications that leverage that connectivity.

But the VRM market as narrowly defined in this document will never be the size of the router market. The technology is just too powerful. Since it creates and operates on virtual entities, one doesn't need a VRM device per LAN connection or WAN link. In fact, several well-placed and highly featured VRM devices can provide a global high availability solution for anywhere between \$200K - \$500K. The product costs are set well below the function value, and competition is driving them even lower.

But I only care about that to the extent that I hope the vendors inventing and honing this capability get rewarded enough to be able to continue their quest. That way, I'll continue to have interesting things to investigate and report on.

2 1999 In Review

2.1 The Big Picture (It's Pretty Exciting)

In the seven months since we published the first two VRM research reports, a lot has happened. A rapid growth of the market, buying guides and product comparisons in the press, IPOs, acquisitions, significant shifts in market positioning, some technology advances, cookie mania, new applications, and even an attempt by a couple of start-ups to commoditize the whole product category.

It was a period that saw F5 Networks, Alteon Websystems, Foundry Networks and RadWare all launch very successful IPOs. Resonate and ArrowPoint are on the IPO track for early 2000. During the same period Intel acquired IPivot for \$400M and Cisco sent signals to the marketplace that they were serious about this product space through a significant change in architectural direction. We also saw Extreme Networks, 3COM, and Cabletron jump into the mix, all aided and abetted by F5 Networks. Finally, the first several \$1M+ end customer deals in this space (at least that I am aware of) went down, with many discussions going on all over the place that could lead to many more.

In other words, in the past 6-8 months many constituencies that make for a hot market “got it” - large customers, large vendors, the trade press, Wall Street, and others. I no longer feel like a lone wolf skulking in the weeds when I spend time thinking about, talking about, or using VRM technology. It's almost become mainstream – at least among the Internet-centric community. The enterprise segment has started to show signs of “getting it,” but I think it take another year or so for significant growth of that segment.

Although small by some standards, we have seen the VRM market more than double in 1999, to a little over \$270M. That represents greater than 100% growth in revenue, and I believe the growth in usage of the technology is even greater than that. This is further evidence that 1999 was a year where various constituencies “got it.”

2.2 Now You Can Start To Measure Success

Some estimated financials:

Company	Symbol	Share Price (3/28/00)	Market CAP (\$M) (3/28/00)	% Rev. VRM	4Q Rev. (\$M)	4Q99 Market Share (Rev.)	1999 Rev. (\$M)	1999 Market Share (Rev.)	Gross Margin	Approx. 1999 Units Sold	Approx. Cust. Base	Cust Base Share	Approx. # Employ. (VRM)
Total Market					\$97.9		\$265.0			23570	8400		
Cisco	CSCO			0.3%	\$21.5	22.0%	\$77.0	29.1%	71%	5500	2400	28.6%	small (vrn)
F5 Networks	FFIV	\$67.75	\$1,433	100%	\$19.2	19.6%	\$44.4	16.8%	74%	3171	1250	14.9%	275
Alteon	ATON	\$82.00	\$3,319	70%	\$12.0	12.3%	\$31.1	11.7%	55%	2704	1000	11.9%	250
Foundry	FDRY	\$143.75	\$15,911	20%	\$11.0	11.2%	\$26.7	10.1%	50%	3141	800	9.5%	
ArrowPoint	ARPT	\$137.50	\$452	100%	\$6.1	6.2%	\$12.4	4.7%	59%	954	225	2.7%	90
RadWare	RDWR	\$38.75	\$566	100%	\$5.0	5.1%	\$14.1	5.3%	85%	1659	700	8.3%	100
IPivot/Intel	INTC				\$4.9	5.0%	\$13.8	5.2%	50%	1725	450	5.4%	120
IBM	IBM				\$4.3	4.4%	\$10.0	3.8%	100%	1667	400	4.8%	small (vrn)
HydraWeb	private			100%	\$4.2	4.3%	\$10.5	4.0%	72%	700	250	3.0%	90
Resonate	private			100%	\$3.9	4.0%	\$9.9	3.7%	100%	707	500	6.0%	175
Other					\$3.5	3.6%	\$9.0	3.4%		900	250	3.0%	
Holontech	private			100%	\$1.7	1.7%	\$4.7	1.8%	70%	392	100	1.2%	80
CoyotePoint	private			100%	\$0.6	0.6%	\$1.4	0.5%	60%	350	75	0.9%	30


Revenues shown are vendor revenues, i.e. the amount of revenue the vendor earned. These numbers under-estimate the total amount of money spent by end users buying VRM solutions because they don't track revenue earned by VARs, distributors, OEMs, etc.

Revenues shown in italicized red represent bottom line losses.

There is a lot of judgment and educated guesswork (i.e. hand waving) in a lot of these numbers:

- For the public companies only shipping VRM-related products (RadWare, F5) it is pretty easy to populate this table with publicly available financials results that they provide.
- For the public companies who do more than VRM (Alteon, Cisco, Intel, Foundry), I have had to use some judgment as to how much of their revenue is related to VRM. Some of that is derived from high-level discussions (accompanied by a wink) from the vendors themselves. But my wink parsing may not be 100% accurate.
- For the private companies who have filed their S-1s (ArrowPoint and Resonate, I had to guess at the amount of revenue in each quarter because while the S-1 documents annual revenue, it does not break it down by quarter.
- For the private companies who have not yet filed their S1-s to go public (e.g. HydraWeb, CoyotePoint), I've had to guess at revenue based on my understanding of the number of customers, ASPs, etc. These numbers are the ones with the greatest potential error factor in them.
- The unit numbers are based on guesses as to the ASP for each vendor. So there are multiple error sources here – the revenue number to start with, and then my guesses as to the product mix shipped and the street price received for each product. Don't take these numbers too seriously. The value of these numbers is to determine a different kind of market share – who is getting the most products out there, creating a base for future growth. Vendors like Foundry and RadWare, who have low ASPs, tend to get short changed a bit when the market share discussions are based on revenue only.

As can be seen from the chart above, Cisco lost market share (again) in 1999, which enabled a large group of alternative vendors to grow and prosper. If you talked to almost any vendor in this space in the 2nd half of 99, they were feeling pretty good about themselves. The market expanded rapidly, and thus they did as well.




Another way to measure success is by year-to-year growth. The chart at left shows the growth from 1998 to 1999 for some of the major vendors. This shows that the growth of several of the vendors has far exceeded the market growth.

Since Cisco was the gorilla of this market in 1998, but lost considerable market share in 1999, it is interesting to compare the various vendor's growth vs. the non-Cisco market growth, which is much greater than the total market growth (percentage-wise). Based on this metric, F5 Networks, Alteon, and Foundry can be seen to be gaining not only against Cisco, but against the other competitors.

There is some inherent unfairness in the chart at left. First of all, neither ArrowPoint nor IPivot are shown. That's because their revenue was so low in 1998 that as far as growth percentage goes in 1999, they'd be off the chart. Secondly, companies that were established and had reasonable revenue in 1998, such as RadWare, Resonate, and Alteon, tend to suffer in terms of percentage growth comparisons because their 1998 baseline is higher than the rest. But still – F5 Networks was established in 1998 – and look at what they did in 1999.

Another way to look success and momentum is by quarter-to-quarter revenue growth in 1999:



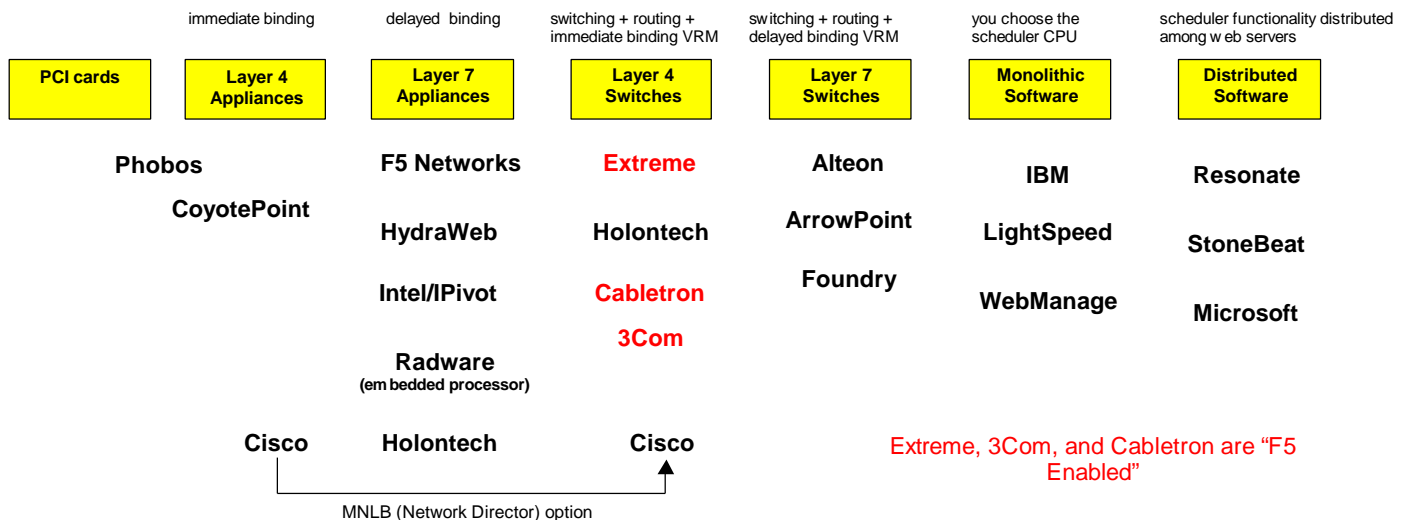
Nobody really faltered in 1999. There is not a dog among the group highlighted above. F5 Networks continued to expand their business and are pretty close to surpassing Cisco as #1 in the market. That will happen in 1Q00. Alteon, Foundry, RadWare, and Resonate all expanded their businesses rapidly. ArrowPoint struggled in the 1st half of the year but regained a lot of momentum in the 2nd half. IPivot (not shown) – a later entrant into the SLB market, made a big splash with their SSL acceleration products and were able to attract enough attention to land Nortel Networks as an OEM customer and then Intel as an acquirer.

As we will discuss in the next section, the slope of the curves on the above figure are as important as the revenue values.

2.3 Summary of Key Technical Trends

2.3.1 A New Taxonomy of Product Types

Life used to be simple. You had your SLB appliances, your SLB switches, and your SLB software. Now, we have a bunch more categories:



What does this all mean? Well, in one sense it means the market has a lot of choices. One can combine VRM with Layer 2/3 switching or not. One can leverage added appliances or not. One can avoid more hardware altogether and overlay the system with software. One can choose Layer 4 or Layer 7 functionality in any form factor.

But in reality, I've come to the conclusion that these aren't really the key factors of the vendor decision matrix. When it comes down to it, VRM is an issue of software value-add capability. There are some 2nd order pros and cons to having that software run on dedicated appliances vs. embedded processors vs. selected server platforms vs. the web servers themselves. But the 1st order issues are the feature set and performance of the VRM software, independent of what processor or processors it is running on where. There are good products of every product type, and they are all rapidly getting both better and more similar.

2.3.2 Cookies for Persistence

One of the most important characteristics of a VRM solution is to be able to match load-balancing policies to the state management approach of the application. For instance, if a user comes in, browses an E-Commerce site, and eventually indicates he wants to buy a book about day trading, where is that information stored? If it is stored at the web server level, then only the web server the user was connected to when he made the selection knows about it. If, a moment later, that same user were load balanced to a different web server, and he went to check his Shopping Cart, he would see nothing in it. If that happened often, the user would be scared off from doing any E-Commerce at the site. So the VRM solution must deliver persistence – the act of keeping the same user connected to the same server for a period of time, consistent with the characteristics of the application.

You might ask, why doesn't the application store the state in the back-end, possibly on the database? If it did, then it wouldn't matter what web server and/or application server the user was connected to at any point in time - they would all have access to the same and most up-to-date information.

The main reason many applications don't operate this way is performance. It would cost a lot for servers to have to do a database look-up every time they wanted to look at or modify a variable. The second reason is simplicity of application design. Both these reasons, combined with inexplicable Microsoft logic, means that if you are running an IIS application, you almost always have state stored at the web or application server level for a period of time. Ultimately, when a transaction is completed or a session is timed out, the data will be stored on the database.

So, for a period of time, the VRM solution must deliver persistence. How do we do that?

One option is to look at the source IP address of the incoming requests and send the same Source IP address to the same server. If you don't see any activity for awhile (you choose a timer value) the VRM system can "forget" the source IP-to-server association so that the same user doesn't get sent to the same server a week later.

The problems with the implementation above are:

- (a) Over time it becomes difficult to load balance effectively. Users may get "stuck" to a server because of occasional interaction with the site, even though no state has been created. If all the users have been "stuck" to a server, you can't redirect traffic to a lightly loaded server and away from a congested one.
- (b) If the user is coming from behind a proxy (as is typical of an AOL user) their Source IP Address may change with each TCP connection, and therefore we don't achieve the "stickiness" we wanted to.

To deal with (b) above, some vendors have added a feature to map all users within a pre-defined Source IP Address range (such as a Class C range) to the same server. This works for stickiness because most proxies allocate addresses within a Class C range. But it just makes issue (a) above even worse. Now you have all users coming from behind a proxy stuck to a server – probably forever because there is a good chance that *someone* in that community will be active at any given time.

So a solution is desired that can be indexed to a unique identifier for a client (as opposed to an IP address, which may or may not be unique). The solution du jour is cookie-based persistence.

Cookies are used for keeping track of settings or data for a particular Web site you have visited, and for identifying you as a unique user to the website that originated the cookie. When your browser requests a page, it sends the settings that apply to that domain along with the request. These settings are stored in a website specific text file on your system.

Cookies can be temporary or permanent. Your browser retains *temporary cookies* while it is operational, and deletes them when it is shut down.

Your browser saves *permanent cookies* as tiny files on your system to maintain settings or data between multiple visits. "Permanent" cookies are actually set to expire at some time in the future (commonly between 30 days and a year from their creation date), and are automatically deleted from your system at that time.

The following shows an example permanent cookie stored on one of our systems when we browsed to Compaq's online store website:

```

ShopperManager%2Fstore
SHOPPERMANAGER%2FSTORE=52NALWB9LKS12G7X00GPBLF689ELAMX1
athome.compaq.com/store
0
2802348032
30050983
1374920960
29243240
*
```

Either temporary or permanent cookies can be used to improve persistence mechanisms related to VRM. For instance, a VRM system supporting a website with a “shopping cart” may use a temporary cookie to identify you and make sure you get sent to the server that remembers your choices from one web page to the next.

When a (cookie-capable) VRM system gets involved, it terminates the TCP connection, waits for the URL requests to come from the user and then performs cookie inspection. The VRM scheduler chooses a server for the user and then records a representation of the user’s cookie within its session state tables, remembering that the user with the inspected cookie is associated with the chosen server. Subsequent requests from the same user will continue to go to the same server, until some timer expires the session state entry.

This persistent binding technique offers more load balancing granularity than IP Source Address based load balancing, because the VRM system is now able to identify a specific *user* to send to a server. It also holds promise for use in identifying specific users or user classes for Preferential Services capabilities.

Issues

There are a number of significant issues with the use of cookies as unique user identifiers:

- *What if clients don’t accept cookies?* If a user has turned off cookie-acceptance on their browser, it is obviously not possible to identify the user by cookie. It’s likely that the website the user is connecting to will complain about its inability to set a user cookie, but if the site has a workaround, what will the VRM scheduler do? It will either have to fallback its persistent policy to one of the other methods, probably Source IP address binding *or it will have to set and recognize values set in the URL as CGI parameters*. Some VRM solutions perform the latter function as an option today.
- *What portion of the cookie does the VRM administrator set persistence on?* Cookies and the specific user-identifiers within them are not human readable, and they are unique from site to site. This opens up a whole set of problems for the VRM administrator. There is no way to simply look at the cookie and determine what to use as the user ID. Plus, the cookie can change structure at any time based on the whim of the website applications programmers. Thus, close coordination between the website programmers and the VRM administrator is required, both for initial setup and during website changes. *One-way around this is for the VRM system to set the cookies itself, using a structure that is easy for the VRM system to understand and parse.*
- *Cookies provided across an SSL session cannot be inspected.* SSL session data is encrypted. Not human readable. Can’t be parsed. Thus, if you must maintain user persistent bindings between both HTTP sessions and HTTP-S sessions, you’ll have to use Source IP address binding or have some method of tracking the HTTP session to the SSL session and then back to HTTP, *or provide the SSL function as a front end to the VRM solution, as IPivot (Intel) does with their eCommerce Accelerator*

boxes, as Phobos does with their SSLxpress boxes, and as Holontech and F5 Networks have announced their intention to do.

2.3.3 Bypass for Show, Delay for Dough

1999 was a year in which many vendors learned a trick to make their performance specifications better: MAC Address Translation and Direct-Path-Return.

Most vendors have always had a MAC Address Translation option, (where the Destination MAC of incoming requests is changed, but nothing is changed at the IP layer or higher) because this mode is needed for Firewall Load Balancing and Cache Re-Direction. But in 1999 it was also added as an option for Server Load Balancing in most products. The advantage of MAC Address Translation is that there is no reason the traffic returning from the server must pass through the scheduler on it's way back to the client. Or if forced to topologically, the scheduler just has to perform a simple Layer 2 bridging function rather than a more intensive Network Address Translation function. The end result is higher performance.

However, there are some things you give up to achieve this performance:

- 1) You must configure loop back interfaces on all the web servers so that they will accept traffic to the Virtual IP address and use the Virtual IP as the source IP address in returned packets.
- 2) The scheduler may not see the returned traffic. Therefore it cannot count it (for better load balancing) or analyze it (for identifying errors).
- 3) Most importantly, this mode of operation disallows all Delayed Binding functions (e.g. URL-based scheduling, SSL Session ID tracking, cookie-based persistence, cookie-based Preferential Services)¹.

Because of the greater performance, many vendors have now gotten into the habit of advertising performance numbers based on the Direct Path Return mode of operation. This can be very misleading because nine times out of ten, that is not the mode of operation the devices are configured to operate in. The real test of a product's performance is how it handles traffic with a heavy load of Delayed Binding, monitoring, security, and management features turned on.

2.3.4 Getting Multi-Site Right?

In 1999, probably about 20% of the local VRM users also use Multi-Site VRM (or Global or Wide Area, or Whatever Your Favorite Term Is). That's up significantly from 1998. But even a lot more people are looking at it, and 40-50% of the vendor decisions being made these days take into account the vendor's multi-site solution, even though that capability may not be deployed until later on.

Most vendors that didn't have a multi-site solution coming into 1999 added one. Those vendors that had a solution improved it. No vendor's solutions are perfect, but a year of usage and on-going development has helped make the solutions very usable. Towards the end of 1999, the state-of-the art for multi-site solutions has started to emerge. Key attributes are:

- The ability to use DNS Re-Direction as the primary method of sending users to selected sites. That is, direct users to sites by responding to the DNS host name resolution process with the IP address of the virtual service at the "best site" determined for that user.

¹ Only one company, Resonate, can perform both Layer 7 services and direct-path-return simultaneously, because of their unique distributed software architecture.

- The ability to take a wide range of factors into account in making the “best site” decision, and to provide a precedence and/or weighting to these factors in making that decision.
 - The source address of the requesting client DNS.
 - The host name of the requesting client DNS.
 - The number of Autonomous System hops between the requesting client LDNS and the various site options.
 - The number of router hops between the requesting client LDNS and the various site options.
 - The packet loss rate between the requesting client LDNS and the various site options.
 - The health and availability of each site.
 - The load as compared to capacity at each of the candidate sites.
 - The WAN latency from the requesting client LDNS to each of the various site options.
 - The status of the site during periods of content and application refreshes, i.e., if a site is being updated, to temporarily take it out of the candidate sites list.
 - Over-riding all the above, a persistency policy so that users connected to site where state is maintained (such as at a local database) remain connected to that site until the state is eliminated (such as via database re-synchronization).
- The ability to layer the rules associated with the above, i.e. use Autonomous Systems hops to narrow the selection to three sites, and then Round Robin (or lowest latency, or whatever) to choose between those sites.
- The ability to re-direct traffic away from a site once it has arrived at a site, to avoid site overload conditions that just occurred.
- The ability to re-measure Autonomous System hops, router hops, and WAN latency to the true user (DNS Re-Direction just measures to the users LDNS, which may or may not be co-located with the user), once that users request arrives at a site, and re-direct future requests to a better site for that user, if the initial decision made by the DNS response system was particularly egregious.
- Support for flexible script-based tuning and customization of the multi-site scheduler policies and parameters, to tune to specific needs. There seem to be as many exceptions and there are rules in this arena.

Up until recently, all of the vendor solutions have used DNS Re-Direction as their main multi-site mechanism. There are some problems with using DNS for multi-site load balancing, the most critical of which are sluggishness of site fail-over due to addresses being cached at both the LDNS and browser levels. Foundry has recently announced a capability that may be a break-through in multi-site function. They call it Global IP, and it potentially allows you to eliminate the DNS Re-Direction process, at least for sites in the same Autonomous System (I discuss the details a bit more in the Foundry section later on in this document). It's new and relatively untested.

2.3.5 Firewall Load Balancing – It's Not Just For Breakfast Any More

One of the fastest growing applications for VRM products in 1999 was firewall load balancing, with some vendors attributing 20-25% of their revenues to this application. This makes sense to me. The firewall is a critical resource that can be a single point-of-failure or a system bottleneck (because its function can be very compute intensive). Sounds kind of like a web server from that perspective, doesn't it?

Firewall load balancing allows multiple firewalls to look and act as a single, more highly scaled “virtual firewall.” At its best, firewall load balancing enables the use of both proxy and transparent firewall services, stateful and non-stateful policies, both application-aware and non-application-aware features, and the ability to flexibly assign different firewalls responsibility for different application groups or user groups.


Although every vendor claims support for firewall load balancing, in some cases, it is limited to support for non-stateful proxy services. In other cases support for transparent, but non-stateful functions is supported. These capabilities are pretty simple. The tricky part is in supporting stateful transparent services because the load balancing system must ensure that both directions of flow for a particular connection pass through the same firewall. To achieve that, most vendors propose a “firewall sandwich” configuration where the Trusted, Untrusted, and DMZ interfaces of the firewall are all connected to a firewall load balancer (FLB) device. For the more sophisticated products, algorithms running in the FLBs ensure that bi-direction traffic flows associated with a connection are mapped to pass through the same firewall.

The “firewall sandwiches” can get pretty expensive because if you make the FLBs redundant, you can need six or more FLBs to create the sandwich (it’s more like a club than a regular sandwich). That’s a lot of equipment and expense if you are just going from one firewall to two or three. Instead, you may choose to just use your firewall vendor’s hot stand-by scheme to increase your reliability (but not necessarily your scalability). One vendor, RadWare, has introduced the concept of “Port Rules” to allow a single FLB box with multiple ports to be logically configured as several boxes to eliminate the need for a multi-box “firewall sandwich.”

Innovations such as these, the FLB requirements for various firewall applications, and how to extend FLB functions to load balance VPN devices, routers, and other network infrastructure devices will be discussed in a future Acuitive VRM Research Report focused on the art and science of “Non-Server Load Balancing.”

2.3.6 Caching and Server Load Balancing: Complementary or Competitive?

Until recently enterprises and ISPs considered using content caching purely as a *defensive* strategy aimed at reducing the amount of bandwidth required to support repetitive content requests over their WAN access links and Internet backbone links, respectively.



However, the use of content caching by enterprise end users has been modest to date. The issues preventing widespread usage have been:

1. Lack of a large enough user population accessing the same content to ensure a sufficient number of cache hits.
2. Management and operations overhead - mostly associated with maintaining proxy caches and proxy addresses on each user’s browser.
3. Minimal WAN bandwidth and therefore cost savings due to the lack of a high number of cache hits and because enterprise networks tend to support a lot of non-cacheable traffic (e.g. e-mail, ftp, ERP, etc.) over their WAN links.
4. Discontented users due to the occasional delivery of stale, out-of-date content.

The use of content caching by ISPs has been somewhat more popular, fueled by the development of transparent caches (that the end users don't need to know about or address), the possibility of a greater number of cache hits due to supporting a larger user base, and the greater economic advantage of saving bandwidth to the ISPs (vs. enterprises).

However, even within ISPs, caching hasn't really taken the world by storm. This is largely due to tension with the original content providers. To achieve a higher probability of getting a cache hit, the ISPs would like to cache content for long periods of time. From the content providers point of view, that creates a delay in the delivery of updated content to their users. If your business depends on the delivery of up-to-date information, such as a news site, such delays are anathema. One way to reduce this problem is to limit the amount of time that caches cache content. But that reduces the number of cache hits and therefore the advantage from the ISPs perspective.

The only way for all the constituents to benefit from caching – the users, the ISPs, and the content providers - is if the ISPs and content providers cooperate. There are a couple of ways that can happen:


1. New protocols are being developed that allow the content providers to communicate to caches to inform them when content has been updated. Thus ISPs don't have to make broad guesses about how long to cache content elements. They can cache slowly changing content for extended periods of time and update fast changing content quickly. Using these protocols, content sources can also identify which content to not cache at all.
2. Content providers can actively work with ISPs or other service providers to actively manage the content resident on caches by pushing their content to the caches and updating the content as needed.

Technically, (1) and (2) are similar, except that (2) doesn't necessarily require the use of emerging standards and protocols. The bigger difference, however, is the business relationship between the content providers and the service providers in each case.

In approach (1), there is no particular business relationship. There is just a technology that helps ensure that cached content is fresh and up-to-date. That does benefit all the constituents. But it benefits the content provider the least because they can't be sure that the ISPs will cache their content. If they are not getting a lot of hits, or if other content providers are getting more hits, then the ISP may not perform caching for them. And whom they cache and whom they don't continually changes, resulting in potentially inconsistent performance for any particular content provider. Thus there remains some conflict between the content providers and the ISPs.

In approach (2), there is a business relationship between the content provider and the caching service provider. Thus content providers can use this as an *offensive* technique, to move their content closer to users globally to ensure better performance, but at the same time be in total control of what content is cached and when it is updated.

Deploying caches globally would be an expensive proposition for any one content provider. This economic reality has led to a new type of service called a Content Delivery



Service (CDS) as delivered by providers such as Digital Island and Akamai, who provision caches globally and then offer usage of it to multiple content providers.

One of the big advantages to a content provider in engaging in such a service is that they can count on it. They are paying for it, so the CDS *will* cache their chosen content. The content provider can take advantage of this by simplifying their core sites. Instead of having large numbers of web servers provisioned at each core site to deliver static content, they can reduce the number of web servers and provisioned WAN bandwidth and focus instead on delivering dynamic applications, E-commerce transactions, and other inherently non-cacheable system components at those sites.

2.3.7 The Myths And Legends Of Security

As I wrote this Industry Update, a rash of distributed denial-of-service attacks were made on several high visibility web sites. I'm sure you read about this.

Many of the VRM vendors discussed in this document immediately saw this as a marketing opportunity and put out press releases or white papers about how their capability could have uniquely either prevented, discouraged, or provided some form of insurance against these attacks.

This very public discussion occurred during a time period where the attacked sites were being very quiet about the nature of the attacks, or were just figuring it out themselves. So the nature of the discussion was generally, "assuming the attack was of a type we understand (e.g. SYN attack, smurf attack, etc.) here is how we could have helped."

And the truth of the matter is, IF the attacks were of the type proposed by the VRM vendors, their solutions could have helped to some significant degree. But, the problem is one of assumption. The nature of these particular attacks for all the affected sites have not been completely characterized. The site I am most familiar with was attacked via a huge uptick in (apparently) legitimate traffic – users who completed the TCP handshake process and then went on to make resource intensive requests of the site.

3 Y2K Market Dynamics

3.1 Key Success Factors For VRM Vendors

The Year 2000 Will Be The Make-Or-Break Year For Vendors In This Industry

Things move fast these days. Markets move through various phases faster than they ever have before. Even though the VRM market seems new and unformed, it has been around since 1996 and I believe that many attributes of a mature, consolidating, and somewhat commodity market will become apparent in 2000, and then even more so in 2001.

Going into 2001, a fewer number of measurably stronger vendors will dominate most of this market. Others will have either faded away, found a niche that they will be forced to be happy with, or gotten acquired (several by Cisco).

For the most part, those that will be strong going into 2001 will be the ones that developed their gathering strength in 1999. In other words, you can pretty accurately identify the winners now.

3.1.1 Just "Be There"

I believe that the market will slightly more than double in 2000, from \$270M to \$580M. However, at least half of the Y2K market will be due to people who already made a vendor decision in 1998 and 1999 making more of the same in Y2K. In the past, we've seen a fair amount of vendor hopping among customers. Every vendor has an anecdote about how they kicked someone else out of an account (and every vendor is the victim of one or more of these anecdotes). But such anecdotes are becoming rarer. The products are more robust, people are learning how to use them more soundly, and the cost of change is going up as people expand their usage of VRM. Therefore, unless some compelling innovation is delivered to the market by a new or existing entrant, roughly \$270M of the \$580M Y2K market is already decided in terms of which vendor's top line it goes to. And I don't see such an innovation on the horizon. Most vendors have roughly the same technology at this point and most new releases are incremental improvements (welcome improvements, but they don't generally disrupt the competitive balance).

But what about the \$310M up for grabs in Y2K that represent buying by new users and customers? It's not a fair race between the vendors for that revenue either.

In 1999, the VRM vendors uniformly reported little competitive pressure, usually in the form of a statement like, "We see Cisco everywhere, but we kick their butt. We don't really see the other guys much." There are elements of macho and head-in-sand with these statements, but actually I believe that for the most part they are true. In 1999 I observed many situations where it took awhile for someone to recognize they had a problem that could be alleviated by VRM, but when they did they wanted to solve it fast and the 1st vendor they became aware of almost always got the deal. So most vendors are having no problem beating Cisco as long as they can get an audience with the prospective buyer. And that's the big issue. This is not a time and technology that lends itself to RFPs, comparison-shopping, bake-offs, etc. People have a VRM itch and they find a VRM scratcher – and fast. The key to success just seems to be "being there." And that means "vibe" and channel development. So when you want to make a bet on which players in this industry will strut into 2001 instead of stumbling into it, pay little attention to new feature, technology, or product announcements. They all have roughly the same technology, and if they don't they will within a few months. Pay more attention to announcements about OEM deals, large VAR

or distributor relationships, cooperative marketing and sales agreements, NSP and colo agreements, and M&A activities. Of course, it is sometimes hard to discern which such announcements are substantive and which are superficial, but somewhere in that pile of information will reveal the success formula of the winners as we go into 2001.

3.1.2 SLB “Break Out” Strategy

Almost every VRM vendor will tell you that they are not just about Server Load Balancing. Sometimes they almost spit when they say “SLB,” like it’s some vile or prehistoric function. I start to remind them that it is SLB that pays their salaries, but then I remember that it really isn’t – when you look at most of their bottom lines it’s usually the venture capitalists and public stockholders that foot the bill.

Anyway, for whatever reason, SLB became a naughty term in 1999, and every vendor goes to great lengths to convince their customers, the press, the analyst community, and I think, themselves, that they are “more than just an SLB company.”

And, in reality, I understand why each vendor needs an SLB breakout strategy. It’s because you can’t survive if you can only sell \$30,000-\$50,000 of complex product to each customer. The R&D, sales, marketing, and support costs more than eat that up and turn that revenue into a red bottom line number. And although I think the intrinsic value of SLB is very high, competition has driven the expectation for the cost of SLB down - to the point where a Resonate, a software company with no COGS, gives it away for free.

OK, so if SLB is free, how do we make any revenue? Here are the tactics du jour:

- Define every feature and function except for the most basic capabilities as “not-SLB” – you have to pay for them. This includes Delayed Binding functions, multi-site capabilities, and management tools.
- Sell related functions that are not specifically SLB – firewalling, bandwidth management, VPN functions, SSL acceleration, content management, caching, usage analysis and reporting, umbrella policy management, routing, switching, etc.

3.1.3 Clear and Unique Positioning

I track 12 vendors in this space, and more are coming on board. For even 6-8 of these to survive they either need to:

- Find a unique niche in the overall market space that allows them to subdivide the market and dominate one aspect.
- Develop a unique message that naturally attracts a significant set of users to them at the market level, across all segments
- Both of the above.

3.2 Positioning Summaries

The table below summarizes each vendor's SLB Break-Out strategy and positioning.

Vendor	Positioning	SLB Break Out Strategy
Cisco	We're the big guys. We've always been there to support you.	Don't need a breakout strategy. SLB is a mite on the flea on the tail of the dog. <i>The best strategy is to be in a position where you don't need one.</i>
F5 Networks	We're the best, the 1 st to everything, and the clear market leader. Our technology will be the de facto standard for internet Traffic and Content Management.	Become the de facto standard for L4 software, in a variety of solutions. Add value added functions (multi-site, management, content distribution, bandwidth mgmt, firewalling) on top of that to become a complete global site management solution. <i>It's working.</i>
Alteon	We're the industry leader in Web Data center switching solutions.	Win more ports in data centers. Get paid for value added functions such as bandwidth management and security services. Define new products for the colo-shared infrastructure. Penetrate the high-end L2/L3 market. <i>Positioning is not completely unique (Foundry, Extreme, ArrowPoint, etc.), which isn't necessarily a bad thing.</i>
Foundry	We've got great routing/switching, and we've got great VRM. If you like us for one, you'll like us for the other.	Don't need one an SLB breakout strategy. Started as a high speed L2/L3 company and are successful there. SLB/VRM just adds to the capability and provides different ways to attract different customers. <i>The best strategy is to be in a position where you don't need one.</i>
Resonate	It's a systems problem, not a networking problem. We're the only vendor that attacks it as a systems problem.	Become an internet software/service company. SLB is free. Offer value-added uplift products, mostly based on an SLB foundation so that customers can build an entire site and multi-site operations system around Resonate products. Drive revenue via Professional Services, Integration Services, and Managed Services. <i>Their very uniqueness ensures them a place in the market. The big issue is whether they can drive systems-oriented decision making from the exception to the norm.</i>
ArrowPoint	It's all about the delivery of content – anywhere and anytime. To do so your control system has to be content-aware.	Don't really need an SLB break out strategy because their positioning from day 1 was at a higher level. <i>Positioning is working. But until mid 1999 execution had slowed ArrowPoint down.</i>
RadWare	We're focused, time tested, and in it for the long haul.	We don't need an SLB breakout strategy. It's a great market and we plan to focus on it and dominate it. <i>Company has good products, good financials, and experience in the space. But they still fight feature battles.</i>
IPivot/Intel	It's all about scalable security and Layer 7 features. You can't have both unless you use us.	Don't need an SLB breakout strategy. Just need efficient integration into the Intel overall strategy and execution process. <i>Watch out for these guys!</i>
Holontech	To make the internet as dependable as the PSTN, you need to use similar management and control techniques.	Leveraging technology sourced from Bell Labs, deliver broad overall fault detection and control capabilities, overlaying the Holontech local and multi-site capabilities. <i>We'll see.</i>
HydraWeb	A quiet company with a track record of industry 1sts. What's the positioning other than being a continual machine of innovation?	
CoyotePoint	SLB is a commodity. You need a sold product that is inexpensive and plug 'n play? We got it.	Breakout strategy is unclear. How do you break out of a market you haven't even broken into yet? <i>The commoditization angle isn't unique. Lot's of vendors have equivalent functionality products for the same price or less.</i>

IBM	SLB is just one of many things you need to deliver E-Commerce. Let's talk about the big picture and the products/services we have to support you.	No need for an SLB breakout strategy because they are already selling at such a higher level. <i>The best strategy is to be in a position where you don't need one.</i>
------------	--	--

But don't forget – execution is more important than the factors scored in the above table.

3.3 Hoover Projections

Now for some real fun. On the next page are my projections for Y2K revenue for the leading vendors of VRM solutions. The middle column entitled “2000 Rev Total” is calculated by summing two factors:

1. Assume that \$270M is spent in Y2K customers who already made a vendor decision in 1998 and 1999, and are just buying more from that vendor. This \$270M is divided up proportionally by the customer base share percentage for each vendor coming into 2000.
2. Assume that customers buying for the 1st time spend an additional \$310M in Y2K. The \$310K is divided up proportionally by the 4Q99 market share numbers.

You can have some fun looking at those numbers, which are inarguable because they are calculated by simple spreadsheet equations. Right?

Well, to have some more fun, I further go on to identify some issues that could result in a vendor either increasing or losing market share compared to the simple calculations above. Such could occur because a vendor either wins/loses business with existing customers or gains/loses market share with new customers.

In general, I give an amplification factor to vendors who seem to have developed increasing momentum in 2H99 and a penalty factor to those who either seem to be losing steam or have bigger risks that other vendors related to channel development, SLB break-out strategy, or other key success factors.

All revenues shown are for VRM only. So for instance, Foundry's revenues for Layer2/3 switches are not counted, nor are Alteon's NICs or the vast majority of the business of Cisco, Intel, and IBM.

My model predicts what I think is a pretty obvious fact – as markets develop and it becomes less about technology, the rich get richer and the poor get (relatively) poorer as the market grows. I think we'll see that in 2000 and then even more so in 2001.

- Based on my model I see the market leaders being F5 Networks, Cisco, Foundry, and Alteon heading into 2001. That's not a very exciting prediction given that they were the leaders going into 2000.
- ArrowPoint should do a good job of trying to catch up with the leaders in 2000, but can't close the gap in one year.
- RadWare will continue on a program of managed, dependable growth.
- Resonate will only have viable businesses going into 2001 if they can establish a clear niche. I think Resonate has and can continue to do that by defining and defending the high end as their space. Resonate will continue to appeal to those who view the problem as a systems problem and prefer a software overlay solution.
- IPivot/Intel will ramp up in 2000 because they will have the company integration behind them and can focus on business.
- Holontech, HydraWeb, CoyotePoint, and new entrants are going to have a difficult time keeping up with the leaders and the strong contenders.

Vendor	4Q99 Market Share (Rev.)	Cust Base Share	2000 Rev. Existing Customer	2000 Rev. New Customer	2000 Rev Total	Upside Scenario	Downside Scenario	MLH Net Up/Dn	2000 Proj. rev. (VRM only)
			\$270.00	\$310.00	\$580.00				\$580.99
Cisco	22.0%	28.6%	\$77.22	\$68.20	\$145.42	Good acquisition early in year; rapid growth of enterprise mkt	Continued loss of installed base to better products; continued slide of new win%	-45.0%	\$79.98
F5 Networks	19.6%	14.9%	\$40.23	\$60.76	\$100.99	SLB break-out works, channels established in 1999 ramp, market leader buying mentality kicks in	Performance becomes issue, colos find a different sol'n for shared infrastructure	15.0%	\$116.14
Alteon	12.3%	11.0%	\$29.70	\$38.13	\$67.83	700 delivery enables SLB break-out; 700 allows penetration of L2/L3	Delay in 700 products delivery; New L4 switch entrants erode uniqueness	12.5%	\$76.31
Foundry	11.2%	9.5%	\$25.65	\$34.72	\$60.37	More capture of VRM revenue with L2/L3 base	L2/L3 market growth stalls; new L4 entrants erode uniqueness	12.5%	\$67.92
Resonate	4.0%	6.0%	\$16.20	\$12.40	\$28.60	Increasing % systems people make VRM decision; Service Level Control SLB break out takes hold	Increasing % network people make VRM decision.	25.0%	\$35.75
Arrow Point	6.2%	2.7%	\$7.29	\$19.22	\$26.51	2H99 momentum builds	Lack of critical mass retards growth	75.0%	\$46.39
RadWare	5.1%	8.3%	\$22.41	\$15.81	\$38.22	Channel builds	Lack of critical mas retards growth	0.0%	\$38.22
IPivot/Intel	5.0%	5.4%	\$14.58	\$15.50	\$30.08	Intel mgmt and channel accelerates velocity	Nortel OEM deal lost; SSL accel. becomes commodity	40.0%	\$42.11
IBM	4.4%	4.8%	\$12.96	\$13.64	\$26.60			-20.0%	\$21.28
Other	3.6%	3.0%	\$8.10	\$11.16	\$19.26	New entrants stumble	F5-enabled vendors take off (counted as part of F5 revenue)	-20.0%	\$15.41
Hydra Web	4.3%	3.0%	\$8.10	\$13.33	\$21.43	Ramps up visibility	Too little channel	10.0%	\$23.57
Holontech	1.7%	1.2%	\$3.24	\$5.27	\$8.51	Channel investment in 2H99 pays off; new mgmt team gels	Too little, too late; continued mgmt churn	50.0%	\$12.77
Coyote Point	0.6%	0.9%	\$2.43	\$1.86	\$4.29	Commoditization is accepted	Too ealry for market to commoditize and/or other alternative commodity choices	20.0%	\$5.15

These revenue projections are for VRM products only.

- Obviously, for Cisco, IBM, and Intel, these projections only represent a tiny part of the company revenue.
- For Foundry, these projections represent about 20-25% of their total revenue.
- For Alteon, these projections represent about 60-75% of their total revenue, depending on how well their NIC business goes in 2000 and how well they penetrate the Layer2/3 switch application space.
- For Resonate, these revenues do not count their Managed Services revenues, which I don't have a good model of prediction for. One could imagine anything from \$5M to \$15M additional service revenue for Resonate.

3.4 My Projections Are Wrong – Here's (Possibly) Why

I am 100% sure that my numbers are wrong in one way or another. I'm just not sure in which ways and how much. Here are some factors that could perturb the model significantly:

3.4.1 More First Time Buyers (Or Less Vendor Loyalty)

If the market isn't split roughly evenly, \$270M related to growth of existing customer bases and \$310M to new customers, then it will shift more towards new customers. Partly, this could be to some customers switching from one vendor to another. We saw a lot of that in 1998 and the 1st half of 1999. Maybe it hasn't settled down as much as I think.

If there is more new buyer \$\$\$ out there, then my model would not be as beneficial to the vendors with the largest installed bases and would be more beneficial to those with the most market share momentum going into 2000. F5 Networks, Alteon, and ArrowPoint would be the most to benefit from that change of assumptions. But I did a sensitivity analysis and unless you assume that all \$500M in 2000 is up for grabs it doesn't change the end conclusions all that much.

3.4.2 F5 Networks Makes A Major Misstep

The only way a significant amount of revenue could be returned to the market pie to be gobbled up by others is if F5 Networks makes a huge misstep in 2000. I can't imagine what that misstep would be. They are doing a good job of developing their products, their business, and in enervating the world with their technology. They seem to me to have too much momentum, backed up by too many brains, to screw up royally. But if they did – the ballgame changes quite a bit.

3.4.3 The Enterprise Market Takes Off

1999 saw some action in the Enterprise (intranet) space, especially as related to mainframe front-ends and firewall load balancing. But still, the majority of the sales were not in the Enterprise space. I assume increased activity is this segment in 2000, but don't expect it to really boom until 2001. If I'm wrong, and it happens in 2000, several consequences will result:

- a) The market will be greater than \$500M – possibly as much as \$600M.
- b) Since a greater part of the overall market will be new buyers than the growth of existing customers, the characteristics discussed in Section 3.4.1 will apply.
- c) Vendors with products and channels most directed towards Enterprise applications – Cisco, F5 Networks (mostly through their switch and systems partners), RadWare (FLB) and Resonate, will benefit a bit more than the rest.

3.4.4 Cisco Gets Their Act Together

This product category shouldn't even exist. Cisco was among the 1st to market and at one point had a huge 70-80% market share. They could have locked everyone else out by investing in the capability and making it just too hard for anyone else to enter. But server load balancers aren't routers, and .com companies don't buy a lot of routers, and the market was/is small, so Cisco pretty much completely starved the Local Director and Distributed Director teams. Not only starved them, but added forced marches (continual re-orgs) to their sentences. Now all the engineers have pretty much escaped.

At some point Cisco is going to see that this functionality is important. Not because of the size of the market per-se, but because of the mindshare. Customers build critical applications around this

technology. If there were enough Cisco people left in the company from the old days, I could have reminded them when a niche technology called routing was getting Cisco people invited to high level customer meetings because it was so critical to building out packet-oriented WANs, even though the spending on routing was dwarfed by spending on SNA components and PBXs and stuff like that. It wasn't about budget, it was about mindshare and the Cisco of old understood that and became the company they are today because of that.

So again, at some point Cisco is going to see that this functionality is important. And once they do, they'll act quickly and decisively. It's too late for them to develop their way back into this market. They are going to have to buy their way back in. Which leads us to the next section.

3.4.5 Colo Implementation Shift

If you go into almost any colo/hoster, you see hundreds of VRM boxes in customer racks. These devices are generally underutilized, take up a lot of rack space and power, and are difficult to manage in large numbers.

It hasn't happened yet, but one could image a VRM product that enabled High Availability and SLB services to be delivered from a switch in the shared infrastructure. Colos would just logically configure this box to provide a service for new customer A, rather than provisioning a box for that function. Potentially, such a box could provide bandwidth management (contract policing), firewalling, and VPN services as well.

If such a box were to be delivered to the market and be bought into by the colos, the nature of the market would shift radically. Literally thousands of boxes that would have been bought for installation into customer racks would not be.

I don't know for sure that this will happen, but it could. The switch vendors or some new vendor seem in the best position to focus on this application and nail it. But I wouldn't put it past the vendors with a close relationship with the colos to recognize this application, like F5 Networks, to figure out a way to serve it as well.

3.4.6 Shift In Buyer Background

Right now most buyers of VRM come from a networking background and view it as a networking problem. This is one reason why roughly 90% of the applications are served by appliances or switches. If for some reason there was a shift towards decision makers coming from the applications/servers area, and looking at it more as a systems decision - then software vendors like Resonate would benefit. As little as a 10% shift towards systems-oriented decisions would have a dramatic impact on Resonate's results, although it would be absorbed and not felt that much by any single hardware vendor. Conversely, a shift towards fewer systems-oriented decision makers and the fortune of Resonate would plummet.

I believe there will be some shift towards an increase in systems-oriented decision makers and decisions, mostly because I think there will be an increase in large scale Enterprise and Enterprise-extension ("click-and-mortar") applications, many of which today are managed by systems management solutions from Tivoli, CA, BMC, MicroMuse and others, whose users are naturally attracted to large scale software deployments.

I have already factored a small shift in this direction into my projections.

3.4.7 Market Consolidation

The one thing that you can almost bet on that will happen in 2000 that will throw my projections off is market consolidation. Everything is prime for a bunch of M&A in this space:

- The big networking companies and the systems companies all need this capability in-house.
- As the rich get richer and the poor get poorer, smaller companies with good technology but who haven't been able to keep pace channel-wise, will be open to acquisition discussions.

The Buyers

Networking Vendors

As described above, I wouldn't be surprised if Cisco made one or even two acquisitions in this area. The Service Provider Business unit needs a product that may be different from what the Enterprise Business Unit needs.

Nortel Networks made an OEM Agreement with IPivot in 1H99. But can they count on Intel to be good supplier to them in the future? Will they end up competing with one another? Even if you could avoid such conflict, Nortel will need to have the capability in-house so they can integrate it with other policy-based functions on a pace and with differentiating features they derive.

If Cisco and Nortel do it, Lucent, Alcatel, Ericsson, Nokia, and Siemens have to at least think about it.

You could also see Extreme making an acquisition in this area, if they feel the technology is strategic and therefore they can't be tethered to F5 Networks forever.

Five or six acquisitions could easily occur just within this category of acquirers.

Systems Vendors

Computer companies have been trying to get into networking for years, with (generally) poor success. Part of the problem is that they have tended to try to get in by marketing products in categories that were already well developed and had established leaders, and from a marketing point of view were difficult to tie into the server and application product stories due to their standards-oriented nature.

If I was a computer systems vendor and I wanted to get into networking, I'd use VRM as a cornerstone of my strategy. It operates at that seam where networks and applications converge. There is room for innovation in integration with servers and applications.

HP, Sun, Compaq and a few others should all be looking at this.

Application Vendors

Application vendors all need to build more scalable solutions that can be distributed across multiple sites. With some additional engineering effort, the software core of most VRM products could become the object broker for a wide variety of packaged applications, application development platforms, etc.

Service Level Management Vendors

Most so-called service-level management products don't provide very much real-time control. Mostly they monitor performance and resource utilization and then notify someone by e-mail or page if there is a problem. Here has been some effort to tie some of these systems into bandwidth management products to help them exert their policies more effectively. But often, the bottleneck is not the network; it's the server

and its related resources. Companies like BMC, Tivoli, CA, Concord, and others should be looking seriously at VRM companies as a way to expand into taking more active control of the flows of application requests into application clusters so that they have a lever to execute real-time service level management.

The Potential Targets

F5 Networks is on a rocket ship. At a \$1.8B price tag, they are actually a bargain relative to value. But probably only Cisco would have the wherewithal to consider such an acquisition. That would be an excellent match. If such an acquisition did occur, the market landscape would change dramatically and forever. But I don't think Cisco feels that much pain in this area yet. So while I think they'll acquire, they won't acquire this big. F5 Networks will be left alone to expand their business and start to dominate this market themselves.

Alteon Websystems is also an expensive buy. Big networking companies probably would be less interested because of product line overlap with Alteon switches and the unattractiveness of the NIC business. Big computer companies, however, would be attracted to both those assets. However, I think Alteon would be viewed as too expensive by the big computer companies, who have gotten burned in the past by poor networking and acquisition strategies. They might do a lower key acquisition should they try it again. I expect Alteon to keep going and growing on their own. They may be an acquirer in order to accelerate their SLB breakout strategy.

Foundry Networks would be a huge acquisition at \$15B or so. You wouldn't spend that much to add VRM to your company. If you were to buy Foundry it would be for their high-end routing/switching, although their good VRM technology would still be a plus. Still – too big. They'll be around as a company next year.

Resonate's solution architecture – distributed software, probably doesn't overlay well with the traditional big networking vendors who have a box-pusher mentality. I could see maybe Cisco being interested to bring Delayed Binding functionality to their Multi-Node Load balancing architecture – catching up with the industry on features in one fell swoop. But other than that, I don't see the networking guys as being the acquirers here. I do feel that a lot of people will be sniffing around Resonate, however. Being a software solution, they could easily be incorporated by any one of a number of systems companies, service level management companies, or application vendors. A lot will depend on how expensive they get after their IPO.

RadWare is a great acquisition target. They have really good technology, great margins, and are profitable. They are growing reasonably fast, but as measured by revenue, not as fast as the rest of the companies they compete with. That means to me that they would benefit from being mated to a company with an existing strong channel that needs the RadWare stuff on their price list. Any of the networking companies or systems companies discussed above would benefit from a RadWare acquisition. At a \$700M market cap, they are a (relative) bargain in this space. I'd be really surprised if RadWare were *not* acquired in Y2K. If I had an extra \$700M floating around, I'd buy'em. But I'm saving my spare change for a major league baseball team.

ArrowPoint could use a home but I don't know if they know it yet. Their technology has stabilized, and their visionary messages are starting to take hold – but other vendors are benefiting as much as they are from content-aware service awareness at this point. At this point of the market maturation it seems to me to be almost impossible for ArrowPoint to catch up on their own, although I think they will probably grow their business on a percentage basis more than any other vendor in 1H00. As a private company they would have been an ideal pick-up for someone. Now that they've gone public, they may be too

expensive for some of the possible suitors, depending on how the financial community treats them in the months following their very successful IPO.

HydraWeb has excellent technology and a good track record of innovative industry firsts. But they just don't have the visibility and channel bandwidth to take advantage of that as the market explodes in 2000 and 2001. They do have good penetration into the Manhattan-oriented financial services sector, which could be as an attractive foot-in-door opportunity for an acquirer.

Holontech made some channel investments in 2H99 that will benefit them in Y2K. They've also put a completely new management team into place. This won't allow them to catch up with the rest of the field in Y2K. But it will allow them to build a stronger business that may be more attractive and valuable as an acquisition.

I expect two or three of RadWare, ArrowPoint, HydraWeb, and Holontech to get acquired in Y2K. These are all \$400M to \$1B+ deals. Resonate may also be a good target for a software or systems company, but the price would be somewhat higher. Smaller players with solid technology like CoyotePoint or Phobos may also get picked up in the \$100M - \$300M range.

4 Vendor Updates

4.1 Alteon

4.1.1 New News since May, 1999

Local Load Balancing Feature Releases

Web OS 6.0 release, Fall 1999

URL-Load Balancing

- Parsing of each GET in an HTTP 1.0 or HTTP 1.1 request and bidding to the “best server” that supports that URL.
- When applied to Cache Re-Direction, allows requests for non-cacheable content (eg CGI scripts and dynamic objects) to bypass the cache.

SSL Session ID Tracking and Persistence

“Cookie” cutting features

- Track cookies for maintaining session persistence.
- Support for looking at cookies (inbound) and applying Preferential Services depending on the traffic class indicated by the cookie.
- Support for looking at cookies (outbound) and setting TOS tag bits to indicate the priority level for the packet for downstream DiffServ-aware routers and switches.

CheckPoint OPSEC Certification for Firewall Load Balancing

Very Successful IPO and Follow-up Secondary Offering

4.1.2 Updated Analysis

No one I know is more deserving of their riches than the three engineer-founders of Alteon. The most fun about observing Alteon these days is watching how they spend their money. One has bought some new sweaters and plans to take some longer vacations in the future (meaning some vacations). A second has probably bought himself a new motorcycle. I hear the third has bought a “Lifestyle of the Rich and Famous” house in the hills of Silicon Valley, complete with racquetball court, movie theatre, bowling alley, and hot and cold running UNIX servers. But none of these people have bought me a beer. Jeez.

Alteon has done a good job of capital allocation, meaning they have been good at balancing investments in their channel with investments in R&D. That, combined with aggressive marketing, has put them into a good market position going into 2000. It is the combination of this market position, the momentum behind it, the capitalization enabled by their successful IPO, and pretty good products, that makes them a strong contender for being one of the dominant vendors in this space going into 2001.

The main issue at Alteon these days is getting their next generation Layer 2-7 switch product line out the door and into the market. Announced last April for delivery last fall, we’re still waiting. Maybe 2Q00, we hear now. The main issue here is internally developed ASICs that used to be a high risk/high reward

game, and now is more of a high risk/modest reward game. Alteon is seeing the risk part of that equation – in terms of product delivery delays. But will they see the reward side – will their products be demonstrably and maintainably better than others who didn't invest in these complex ASIC developments and instead chose to leverage merchant silicon (network processors, switch fabrics, etc.)? From the point of view of a stand-alone company, it's not clear how that will play out. But in terms of being an acquisition target, it gives Alteon an extremely valuable technology asset.

The new Alteon switches are critical to their "SLB break-out" strategy. Alteon is in a situation now where they have to spend a lot of time explaining the value of their technology in a complex application space, and helping users to deploy and configure their products, all for customer wins which generally result in the sale of 3-5 boxes, representing maybe \$50,000 of revenue. That's a lot of work for a fairly small amount of revenue. The problem is that Alteon boxes are too powerful – they have good performance, you can do all forms of load balancing from them, including multi-site load balancing -and they were under priced. That last fact is due to the fact that initially Alteon just viewed VRM as a feature to differentiate their Layer 2 switches, rather than a product category with a unique value in and of itself.

If Alteon continues to earn \$50,000 per customer, they can grow revenue forever and never turn profitable.

So to break out of this, to leverage the channel and customer base they have developed, and to earn more revenue per customer, Alteon needs the new switches. These switches may do five things for Alteon:

- They will eliminate some shortcoming relative to limited capacity in the existing products. The present products exhibit blazing speeds and lots of features, but only if the usage of those features is modest in terms of the number of connections, URL parsing rules, etc. This is due to limited memory. The next generation products will alleviate the memory limitation to a large extent, so that you can actually get the benefit of some of the Layer 7 features in real-world situations.
- They will provide a platform to support more features such as firewalling, bandwidth management, and encryption/authentication that potentially can be packaged and priced-to-value separately so that Alteon gains revenue when end customers buy these features. Potentially Alteon could also do this for Layer 4 and Layer 7 SLB features, cache re-direction features, and firewall load balancing features so that they can maintain both a competitive price for applications that don't require these features and fair value for applications that do.
- They will allow Alteon to sell more into service provider infrastructures that want a more robust and modular physical platform.
- They will allow Alteon to sell more ports per customer win. Right now, to support a lot of servers, you have to surround the Alteon box with a lot of Layer 2 switches or hubs. With the new switches Alteon's position will be that you should just plug them directly into Alteon ports.
- They may allow Alteon to penetrate the high end Layer 2/3-switch market. So far, Alteon has not penetrated this market at all. If you buy an Alteon switch it is for VRM, not routing or switching. This is as opposed to Foundry, who has penetrated that market (in fact, that is where they got started) and therefore has 4x more revenue and 4x higher market capitalization.

My own belief is that the first four bullets will definitely come true when Alteon delivers on their new platform. The 5th bullet will also come true to some small degree, but I think it is too late for Alteon to catch up to Foundry, Extreme, and all the big systems vendors in this area. Alteon, like ArrowPoint, should focus on the VRM and related high value aspects of their product line. In doing so, they will also sell some additional boxes for non-VRM functions that are required on the periphery of the VRM applications, to their VRM customers.

In spite of the delay of the delivery of the products that will bring all the wonderful rewards listed above, Alteon is doing very well in the marketplace in terms of revenue. But they are not yet profitable. To become profitable, if they kept their recurring costs fixed, they would need to increase their quarterly revenue by about \$8M. But they are not keeping their recurring costs fixed. They are rapidly hiring and growing, adding about \$1.5M of cost per quarter. So to be profitable they have the following quarterly revenue targets:

	4Q00	1Q00	2Q00	3Q00	4Q00
Operating Loss	\$4.0M				
Increased Recurring Cost	Baseline	\$1.5M	\$3.0M	\$4.5M	\$6.0M
Revenue (total, not just VRM)	\$17.1M				
Break-Even Revenue		\$28M not likely	\$31M possible	\$34M almost for sure	\$37M

With the greater revenue-per-customer attainable with the new switch product line, Alteon should be able to turn strong revenue growth into something that starts to show black on the bottom line in either 2Q00 or 3Q00.

4.2 ArrowPoint

4.2.1 New News Since May, 1999

New Feature Releases

All of ArrowPoint's new product releases in the last 6-8 months have been oriented towards software enhancements to their existing CS-100 and CS-800 platforms.

"Flash Crowd" Support, 9/99

- Switches detect a high and rapidly increasing volume of requests for particular content.
- The switch replicates content on designated overflow caches or servers, and then starts to balance the incoming load among the original servers as well as the new ones.
- Part of the capability is the ability to state and replicate content among a large number of local and remote servers and caches.
- *For local SLB applications, this feature is most oriented towards cases where content is stored on local server and cache hard drives. In cases where the content is shared on the back end via file systems or databases, no content replication is needed, although the flash crowd feature could still be useful to re-assign the number of front-end web servers allocated to serve specific content. Also, the replication feature could be useful to configure remote sites to serve content it previously was not assigned to serve.*

Reverse Proxy Cache (RPC) switching, 9/99

- An extension to ArrowPoint's already existing proxy and transparent cache re-direction features.
- The RPC switching allows ArrowPoint to front end Reverse Proxy Caches, most likely situated at a Content Service Delivery site (e.g. Digital River), to
 - (a) Improve cache performance by providing a bypass path for non-cacheable traffic (via NAT peering).
 - (b) Improve RPC reliability and scalability by enabling clustering of RPCs to create a larger virtual RPC.

Firewall Load Balancing, 10/99

- Firewall load balancing
- Some native firewall services, including "ping flood" attack protection, per flow content request filtering, and high speed access list-based filtering. *These features go well beyond firewall load balancing. Depending on your needs, you could view these as "firewall elimination" features.*

E-commerce Content Services Suite, 11/99

- Cookie Switching – Full URL parsing to discover embedded cookies anywhere in the HTTP header. This function can then be leveraged to track permanent or temporary cookies in order to either ensure shopping cart session "stickiness" or to provide specific users access to specific content at predefined service levels.
- Encryption Key Caching – Allows the server to re-use SSL Session IDs, which reduces CPU overhead for negotiating new shared secrets.
- SSL Session ID Time-Out Resolution – If a web server times out an SSL session after a period of inactivity as can happen say, when a user is interrupted while providing user and credit card information, the ArrowPoint switches save the Session ID for a long period of time, so that when the

user returns to activity, the next packet is sent to the same server they were interacting with before being interrupted. The server will establish a new SSL Session ID – but no application state information will have been lost and the user can just pick up and complete the transaction.

- Hybrid Transaction Binding – allows Shopping Cart transactions to be made persistent even if the source IP address changes as the session passes from the HTTP mode to the SSL mode. This requires some help on the web server that serves the page that executes SSL-protected transactions. When the page is hit, it must return an absolute HREF indexed to the V_IP at a pre-determined TCP port that is unique for the server. When the ArrowPoint switch sees the SSL Hello coming from the browser, addressed to that V_IP and destination TCP port, it sends the request along to that server. *This is unique and clever.*

XML-Based API

Allows server applications to change services and content rules. Developed for the WebSpective application, but not limited in use to WebSpective.

Hardware Enhancements

Upgraded Processors (January 2000)

Upgraded processors on the CS-800 policy engine, resulting in a 3.5x increase in Layer 7 processing power.

New Cards (4Q99)

4 Port Gigabit Ethernet card for the CS-800

New Partnerships

WebSpective

WebSpective focuses on content management and server monitoring. By providing interfaces between the two products, the ArrowPoint switches can be automatically notified when content is updated or added to various target servers. In addition, server faults or overload conditions detected by WebSpective can be communicated to the ArrowPoint switches so that they can direct traffic away from those servers.

Alcatel OEM Deal

Novell

For ICS

Exodus

As infrastructure for their ReadyCache service and as a reseller.

Global Center

As infrastructure for their mail service and as a reseller.

Very Successful IPO

3/31/00

4.2.2 Updated Analysis

If I had founded a VRM company, it would have been like ArrowPoint.

I have always found ArrowPoint to be the most visionary of all the vendors. From day one, they promoted the concept of “Content-Aware” features and services, at a time many of us were still trying to figure out where addresses were hidden in higher-level layers so that we could perform NAT properly. They don’t need a “SLB break-out” strategy because their whole positioning and messaging from the beginning was at a higher level. They were also the first vendor to find a good balance between software, hardware, and hardware-assisted software for implementing complex VRM functions, and were the first to see that the future was Network Processors, not internal complex ASIC design. That vision has allowed them to focus more resources on value-added software development, which will help them in the long run. As examples, just in the last six months, they have invented and delivered on the concepts of reverse proxy cache switching, hot flash initiated content replication, encryption key caching, and hybrid transaction binding, as well as being among the leaders in applying cookie-aware functions.

But all that vision has had a negative rebound effect as well. The complex features are difficult to explain – both when and why you need them, and how they work. ArrowPoint has had to do a lot of market education. And often – people just don’t need those features. They have more basic requirements. It has also been difficult for them rapidly develop, test, and deliver these complex features. That process has been exacerbated by the fact that there are more variables when you are trying to write software to Network Processors – which themselves have quirks and bugs that you are the first to discover if you use them in new and imaginative ways.

So, for a long time, ArrowPoint’s marketing was ahead of the headlights of the product. This has resulted in some disappointed early customers and at least one disappointed analyst who wanted them to do well because he tends to favor the visionaries.

I believe this all changed around mid-year, 1999, however. Around June or so, I started getting much better reports about the ArrowPoint product stability and performance from a variety of sources. That situation was nicely timed with the fact that for a long time (and even now), ArrowPoint was the only vendor to support VRM on a modular and hardened chassis product. At the same time, some excellent technical marketing on their part has resulted in more general visibility and education about the need for content and cookie-aware features. So in that sense, the world is catching up to ArrowPoint. The time that has gone by has allowed them to make their products a bit more robust and on the marketing side, to repeat the same story over and over again so that people get it.

In their recent S-1 filing, ArrowPoint reported \$12.4M for 1999. I believe most of that was in 2H99. They have moved from less than 1% market share to over 6 % (as measured by 4Q99 revenues). So their slope is great, but their starting point is low. They only have a about 200+ customers, so they are not starting 2000 with nearly the base some of their competitors have. The big question is whether they have enough critical mass at this point to sustain and possibly even accelerate their growth, so that they can be one of the leading players going into 2001.

There are two possible scenarios for ArrowPoint. One is that they slingshot into being one of top few vendors in this space because they have innovative features which are well marketed, and they pulled it together product-wise at just the right time. That’s “the vision wins out” scenario”. The other scenario is that they’ll find that at this point technology doesn’t win - presence and channel does -and they just got there too late. That’s the “execution wins out” scenario, in which they will become a good example for some future technology marketing book about the importance of focusing early, being among the first to market, and then building the business out step-by-step from there.

For almost any other vendor coming into 2000 with their market position, I'd tag them as a "too little, too late company" and bet on the 2nd scenario. But ArrowPoint has shown me enough in 2H99 that I'd place a bet on them. If they don't slingshot into being a market leader, I at least think they'll be high up on a lot of lists come M&A due diligence time. Of course, that won't happen if they get into the awkward position of being *too* highly valued after their IPO. But let's not cry for them in that situation.

4.3 Cisco

4.3.1 New News Since May, 1999

For several years, Cisco has talked about adding VRM to IOS and supporting it on their high volume switch and router products. They finally did it... Sort of.


Early in the summer of 1999, Cisco started talking about their Multi-Node Load Balancing (MNLB) architecture.

Basically, what MNLB does is eliminate objections related to having critical networked data pass through a PC platform of dubious reliability and scalability. The architecture does this by separating the packet-by-packet processing functions (changing addresses, changing port numbers, recalculating FCS) from the session-by-session (binding) and other background control functions (evaluating server/service health and load).

The element that performs the session-by-session control functions and background tasks is called the Services Manager. Right now, the Services Manager function runs on a PC, just like Local Director. In fact, you can think of it as just a mode-of-operation choice for Local Director. Cisco has stated that they plan to physically integrate this function onto blades that plug into their modular switches and routers. I believe this blade is nothing more than a PC-on-a-card.

Services Manager

- Software runs on LocalDirector
- ContentFlow Flow Management Agent
- Makes load-balancing decisions
- Uses MNLB to instruct Forwarding Agents of correct server destination
- Uses server feedback protocol to maintain server capacity and application availability info




The packet-by-packet processing is performed on the switches and routers themselves, using the latest IOS release which includes the functionality Cisco calls the Forwarding Agent. Multiple routers or switches can be used to spread the load for both reliability and scalability reasons. The routers perform a table look-up in reference to incoming TCP SYN packets directed to a V_IP. If there is already an entry in the Table, the router references information such as the Real Server address, performs NAT, and sends the packet along. If there is not an entry, the router sends the packet to the Services Manager, which determines the “best available server” at that time, configures the forwarding tables in each of the routers, and sends the packet back to one of the routers. All other packets associated with the session bypass the Services Manager.

Forwarding Agent

- Cisco IOS router and switch software
- ContentFlow Flow Delivery Agent
- Uses MNLB to communicate with Services Manager
- Sends connection requests to Services Manager
- Receives server destination from Services Manager
- Forwards data to chosen server

Workload Agents

- Runs on either server platforms or management consoles
- Maintains information on server work capacity and application availability
- Communicates with Services Manager using server feedback protocol
- For IBM OS/390 systems delivers OS/390 Workload Manager data



Furthermore, although not technically coupled to the function distribution described above, Cisco has created a protocol (Distributed Feedback Protocol, DFP) to enable integration with 3rd party intelligence such as application monitors, database monitors, etc. They call such server-resident entities “Workload Agents.” They have indicated a commitment to driving a multi-vendor standard for this protocol.

4.3.2 Updated Analysis

Cisco is the Proteon of the server load balancing industry. They had it. Then they lost it. At one point Cisco had a 70-80% market share in this space. Their market share has been rapidly eroding over the past two years as smaller, more focused companies have built better products and have marketed them aggressively. Today, due to the strength of the Cisco marketing and sales machines, they still sell a bunch of VRM boxes, but mostly to customers who don't have the time to shop around and just buy what Cisco tells them to.

Why has this happened? Mostly, because it doesn't matter to Cisco that it has. From their perspective, this is a tiny market. Also, many of the buyers and users of VRM are content providers and web sites who are not strategic customers for Cisco – how many routers does a web site need? So Cisco has put this capability on a low priority, so low in fact that only a handful of engineers maintained the products for a couple of years (and most of those have recently left Cisco) and little field training was provided.

Now, the roll-out of the Multi-Node Load Balancing (MNLB) architecture (discussed in the section above) this summer does indicate that finally the IOS organization has put a little bit of energy into VRM. That architecture does eliminate some big warts associated with Local Director. It removes LD from the packet path, replacing it with one or more switches or routers, which should improve both performance and system reliability. LD becomes more of a system controller off to the side. That function will eventually be moved to an embedded processor on a card in the modular products. In addition, the new architecture creates pathways for server-resident software agents and resource managers to report health, load, and status information back to LD so that it can make more informed decisions.

These architectural changes are very good, but to move to the new architecture, Cisco had to put a hold on new feature development. They were already behind the market by a year or so even before that. So Cisco is primarily still a Layer 4 device with no URL, cookie, or application level functionality.

I personally believe that eventually Cisco gets everything right. I think that they will see pressure in Y2K to fix their product line in this area, both from service provider partners and from a growing amount of enterprise activity. Once they come to the conclusion that this capability has customer mindshare significance out of proportion to the market size, they will move quickly to rectify the weakness in their product line. They will do this through acquisition, maybe even two. No vendor is off the potential list for Cisco. Well, maybe they wouldn't acquire Intel in order to get the IPivot product, and they've already gotten all they want out of IBM (and then some). But the rest are all fair game.

4.4 Coyote Point

4.4.1 New News Since May, 1999

Hardware Changes

The main things are new packaging into 2U high boxes and cost reductions of the E-350 (to \$5,995) and the E-450 (to \$9,995).

Release 4.0

- Active Content Verification – tests all servers in the cluster to verify that content can be requested and returned.

4.4.2 Updated Analysis

The CoyotePoint people are good people and the CoyotePoint products are solid. They have all the basic bells and whistles you'd want for both local and multi-site load balancing. But they don't have some of the more advanced (and admittedly sometimes esoteric) features of the products offered by vendors who have been around longer.

The big issue with CoyotePoint is that they came into the market after a lot of other players were already established. They have not yet established a unique niche or approach that separates them from the rest. That makes it difficult for them to punch through and establish a foothold in the market.

They can say accurately "our products are solid, inexpensive, and easy to install." But that's a commoditization message and the VRM market is not yet ready to be commoditized. Most decisions are still based on value and vendor support more than price. Also, there are other products out there that are in the same price range, or even lower (check out Resonate's free SLB, for example). And when the VRM market does commoditize, the established players with whom we developed channels are the ones who benefit, not up-starts. So the positioning is both tenuous and non-unique.

CoyotePoint doesn't only need an SLB break-out strategy, they need an SLB *break-in* strategy.

4.5 F5 Networks

4.5.1 New News Since May '99

Local Load Balancing Feature Releases

Big/IP Release 2.1

- Started shipping in August 1999
- Cookie-based persistence mode. The 1st in the industry to do so. (Followed quickly by several others – see an upcoming research Report on this multi-faceted issue) Cookies can be set by the Big/IP, or set by the server and tracked by Big/IP.
- Destination Address Affinity. This is most useful when load balancing caches. It's function is to re-direct requests for the same content back to the same cache, resulting in a higher number of cache hits and more efficient use of cache memory.
- Networked-based fail over. Eliminates the requirement for an RS-232 connection between redundant boxes. Keep-Alives and configuration information are sent from box to box over the network. This allows the boxes to be physically separated, so that they can be put on different power grids, for example. The RS-232 connection is retained as an option for faster fail over.
- Connection state fail over (F5 calls this mirroring). This feature maintains synchronisms between the forwarding tables on the active scheduler and the back-up scheduler. This means that connections in process will not be dropped upon scheduler fail over. This is real important if one is halfway through a 90-minute download, for instance. It is less important for HTTP access to small records.
- Extended Content Verification features enhanced to include firewall testing, cache servers, and proxy servers.
- Support for pre-defined Extended Application Verification tests for FTP, NNTP, SMTP, POP3, and MS SQL. F5 previously had the ability to check these services but it took more work to configure and set up the tests,
- 802.1q VLAN Trunking. Allows the Big/IP to tag packets with a specific VLAN ID so that downstream switches can apply VLAN forwarding rules to segregate traffic into user groups.
- NPath™ Technology. This is F5's term for supporting MAC Address Translation, which allows return traffic from the server to bypass the Big/IP on it's way back to the client. This can significantly improve performance because it takes a big load off the scheduler, but it has the disadvantages of needing to configure loop-backs on the servers, eliminating the ability to observe and meter the return traffic to aid load balancing, and the mode can not be chosen if you want to track cookies, SSL Session IDs, URLs, or anything that requires Delayed Binding. Largely, it's just a way to be able to brag about higher performance numbers.
- Secure Network Address translation. This is a traditional firewall feature.

Big/IP Release 3.0

- Shipping in January, 2000
- Active-Active Controllers. F5 does not use the term the same way Acuitive does. We would call this feature "Active Standby". This is the characteristic where each scheduler can be active for different services, but act as back-ups for one another, so that if one fails the remaining scheduler takes over operation for all services. Previous to this release, the F5 back-up scheduler didn't hand any actual traffic at all while the primary was healthy. *Note – we reserve the term Active-Active for the case where two or more schedulers can simultaneously be supporting the same service, using the same set of real servers. This is not a hugely important characteristic unless except in situations where the scheduler can potentially be a system bottleneck, then it is vastly important.*

- URL-based scheduling. Requires turning on Delayed Binding. The performance in this mode is not specified. (It's low).
- Support for bi-directional load balancing of firewalls. When creating a "firewall sandwich" for transparent stateful firewalls, it is important that all traffic from a session in both directions pass through the same firewall. This feature adds support for that application.
- Hash-based cookie persistence. To perform exact matches on cookies requires a lot of processing and slows down system operation. If all you are trying to do is ensure persistence (as opposed to providing preferential service to specific users), you don't have to perform exact matches (which was the approach used in release 2.1). Hashing and ensuring persistence based on the hashed value is fast, efficient, and provides just as good persistence, with only a small degradation of load balancing granularity (becomes small groups of users will get hashed to the same value and sent to the same server, and you can't separate them out if that server gets overloaded). This is a very useful feature.

Multi-Site Load Balancing (3DNS) Feature Releases

3DNS Release 2.0

- Shipping in January, 2000
- Support for sites that don't have Big/IP. This is achieved in two ways:
 - Health checking at the V_IP level if there is another vendor's VRM product on the site. No load or delay information is obtained and taken into account – just availability.
 - Direct interrogation of either stand-alone servers or servers behind a 3rd party VRM device using SNMP probing. Supported if the server has UCD snmp, Solstice Enterprise, or NT/4.0 SNMP agents. If there is a Big/IP at the local site, the Big/IP can do the polling for the 3DNS, saving WAN bandwidth and preventing network errors from causing outages of unreliable transported SNMP/UDP packets.
- Production rules that allow administrators to change 3DNS configurations and rules at specific times of the day or under specific conditions. 3DNS supports a scripting tool to aid in the definition of these rules.
- Integration of NameSurfer™ from DataFellows, which provides a spellchecker function for DNS administration
- Support for router hops as a load balancing metrics. This brings the list of options to the following:
 - Round Robin
 - Ratio
 - Least Connections
 - Random
 - Round Trip Time
 - Completion rate
 - Big/IP Packet rate
 - Global availability
 - Topology
 - Dynamic Ratio
- The last one (Dynamic Ratio) is new to this release. The intent is to solve a previous problem of sending all traffic to one single site (because it has the lowest load at the time) and overloading that site (and then sending everything to the new site which has the best performance until it is overloaded, etc.). Dynamic ratio allows you to set policies such as – send X% of the traffic to the best site, send Y% to the second best site, send Z% to the third best, etc. That way traffic is spread across multiple sites even one site is showing a little better performance at one particular time. The ration rules can be set to change dynamically based on various metrics or thresholds you set.

Channels/Partnerships

In 1999, F5 Networks moved from being a technology leader to being a business leader by establishing some very compelling technology/OEM relationships and by fortifying their relationships with a few key colo-hosters, who have proven to be the best channel for VRM technology to date. F5 has people on site at Exodus for example, to help Exodus plan, deliver, and support services executed with F5 Network technology.

Switch Vendors

On the partnership side, F5 Networks has announced OEM programs with several Layer 3 switch vendors:

- Cabletron
- Extreme
- 3Com
- Top Layer

The details are slightly different in each case, but in general these relationships have two aspects:

- 1) Licensing of F5 Networks Layer 4 local SLB code and integration into various switch platforms marketed by each vendor.
- 2) Reselling by the Layer 3 switch vendor of the rest of F5 Networks product line (for Layer 7 SLB, multi-site, enhanced management and content management). This opens up huge channels for F5.

Systems Vendors

F5 has announced a reseller relationship with HP. It's not clear if this relationship is any deeper than ones established by other VRM vendors with HP in the past. But I do believe this indicates a strategy expansion for F5 – to get closer to some of the systems vendors whose servers are after all, what SLB is all about.

As evidence of that, Dell and F5 recently made an announcement where Dell will enhance their Dell.com in-a-rack solution by porting F5 Big/IP code onto one of their appliance platforms, and sell it along side their PowerApp web and cache servers. This relationship, more so than the HP one, exemplifies the F5 strategy.

Very Successful IPO

Year 2000 Statement of Direction

- Next generation Big/IP.
 - Vastly improved scalability via the use of multiple PCI network-processor based co-processors for all packet-by- packet repetitive L4-L7 functions, leaving the main CPU free to support control functions and exception processing.
 - SSL acceleration via hardware assist on PCI co-processor cards. These cards will start shipping in April 2000.
 - *It appears I have to take back everything I have ever said or written about the performance limitations of appliance-based VRM products. By leveraging PCI cards for various functions and integrating them efficiently into the main motherboard system, it appears that **augmented appliances** such as F5's can achieve very high levels of performance, yet maintain their original value of rapid deployment of features (simple*

development target environment) and low cost (use of high volume industry-standard components and subsystems).

- ReServer – designed to support distributed global content delivery. Really an integration of existing F5 capabilities, with some new enhancements, to help every service provider become an Akamai.
- BIG/IS – A rack-and-stack complete solution for Internet Traffic Management, combining F5 products with best-of-breed bandwidth managers, web servers, cache servers, firewalls, etc. The recent announcement with Dell is an example of the rollout of this strategy.

4.5.2 Updated Analysis

F5 Networks is the Cisco of the server load-balancing industry. They remind me of Cisco back in 1986 through 1994 or so as Cisco went through the process of

- (a) Continually creating a technological separation from their competitors by being focused on software-oriented value-add
- (b) Locking up channels to customers through a series of strategic OEM deals, technology licenses and key distribution outlets
- (c) Training the world on their technology so that lack of user knowledge was not an impediment to market growth
- (d) Expanding the market for their products by finding new applications for the core technology and embracing adjacent technologies and applications
- (e) Continually articulating their longer term vision in understandable terms while at the same time delivering a stream of product upgrades
- (f) Supporting higher speed hardware platforms to house their value added software at the last possible moment – to once again block entry points into the market by competition
- (g) Hiring all the talent they can to create a brain drain for the rest of the industry

This list represents a general purpose menu-for-success for any systems company developing a market for a new high value and rapidly changing technology. In other words, it applies to VRM. And more than any other VRM vendor, F5 is following this menu. But time is compressed these days, so F5 Networks deserves praise for sequencing these actions in a more concentrated fashion.

F5 Networks has a long series of industry 1sts to brag about. Early on, these tended to be feature-oriented 1sts, but now the 1sts are being intermixed with important little things like 1st to IPO and 1st to profitability.

When I meet with the F5 people now, the talk isn't about features and specs. It's about overall traffic and content management solutions, strategic relationships, international strategies, vertical market initiatives, information transfer scalability via new training and customer support initiatives, and their increasingly important role in the overall techno-business fabric of the world. But that doesn't mean their technology is lagging behind. They had the most to talk about in their "What's New" section.

In other words, F5 Network is a grown up company doing grown up things very well. They are operating smoothly on all facets of the success equation and as a result turned profitable in 3Q99 and expanded that profitability on 4Q99. Temporary Cisco unit shipment numbers aside, **F5 is the industry leader** and will continue to be so in the foreseeable future.

Oh – and another F5 Networks characteristic that reminds me of Cisco – their competition *hates* them. That's generally a very good sign.

The only things I can see tripping F5 Networks up are:

- Performance issues related to supporting complex Delayed Binding features on a single processor PC-based platform. *But I believe the next generation Big/IP recently announced will completely eliminate this potential issue.*
- Inability to drive revenue outside of the SLB space, so after they establish their (large) market share in SLB, they get growth bounded by the growth of that (relatively) small market segment.

The last bullet is not entirely in F5s power to manage. But it is partly. And I see them doing everything conceivable to define for themselves a much expanded market in the future. If I had to bet a six-pack, I'd bet on their success in that area as well.

4.5.3 Prospects

It's a little too early to say that F5 has won. But they are two sets up in a five set match. By observing the pattern recognition of their deals, you can see that partners of various sorts are lining up to work with them. This, along with a compounding of reference-based wins, is what results in amplified spacing between industry leaders and also-rans as a market matures. F5 is starting to get the benefits that a market awards to the leader.

The other vendors better start stealing their momentum fast or else next year I'll be saying that F5 Networks has won. The only questions at that time will be how big a piece of the pie they will control and if they will be able to expand the pie by finding new applications for their technology and tackling adjacent markets. The other vendors will be left competing for #2 and #3. Anything after that doesn't count.

4.6 Foundry

4.6.1 New News since May, 1999

Firewall Load Balancing, 9/99

ServerIronXL, 11/99

- Upgraded ServerIron platform with a 400 MHz PowerPC CPU.
- Supports up to 80,000 TCP connections/sec (immediate binding mode).

MAC Address Translation (MAT) and Direct Path Return, 11/99

- This feature is dubbed “SwitchBack” by Foundry.

BigServerIron, 11/99

- Initial support for a SLB on Foundry’s high-end modular boxes. The first release supports immediate binding with MAT and Direct Path return only.
- In the future, Foundry intends to support Delayed Binding and Layer 7 features on the platform via software. Ultimately, a plug-in card with hardware assist for Layer 7 features will be supported.

“Half an ACK Whack”

- Voted “best named feature of 1999” by Mark Hoover.
- It’s a SYN attack security feature.
- Works by maintaining a circular list of uncompleted connections. When a new connection comes in, you replace the oldest one. Uncompleted connections also get timed out if the timer expires before new connections bump it off the list.
- Basically, this creates a SYN arrival rate limiter.

URL-Based Scheduling, 1/00

- Shipped as part of IronWare 6.0.
- Parsing rules based on prefix, suffix, or pattern matching.
- Presently supported on ServerIron only.

Cookie and SSL Session ID Based Scheduling, 1/00

- Shipped as part of Ironware 6.0.
- Presently supported on ServerIron only.
- Can match on a server-initiated cookie or can create a configurable cookie associated with the server a request is forwarded to.

Global Load Balancing

- Installs in front of existing DNSs (i.e. does not become the Authoritative Name Server itself, but intercepts packets returned from ANSs and modifies them).
- DNS Re-Direction is the main mode of operation, but also supports two other modes:
 - HTTP Re-Direction. Generally used as a method of last resort, if a site’s server cluster is overloaded and traffic should be shed to another site.
 - “Global IP.” A method for two sites to advertise the same V_IP via Host Route injection. The routers of the ISP just view this as a single endpoint that can be accessed via two different paths. The interior routing protocols within the ISP make decisions on the best

path to that end point, based upon the ingress point of the client. Traffic is naturally spread (but not necessarily “balanced”) between the two sites. The big advantage is site fail over. If one site goes down, the routing protocols will naturally start sending all packets to the remaining healthy site (assuming there are only two sites). This is a unique capability supported by Foundry, although IPivot had something similar but it wasn’t efficient because you had to advertise the entire subnet the V_IP was on and not just the specific V_IP, as you can in Foundry’s case. This is an advantage of being a real router. The potential disadvantage is that if the interior routing protocols change their view of “best path,” possibly due to interior links going down or being congested, traffic associated with a persistent session could be routed to site A for the first half of the session and then all of a sudden to site B. So, in practice, you only use this feature for site fail over and not for load balancing.

- “Best site” choice can be made by mapping globally assigned IP address ranges, site load, site response time, and WAN latency. In the future, BGP router metrics will be added as an option.
- Foundry’s implementation provides support for “layered rules,” i.e. reduce my list of target sites to the ones in the same continent via IP address range mapping, then reduce those choices by the sites that are operating at greater than 80% utilization, then chose the best site from the remaining choices via WAN latency.

Foundry’s Global features are well thought out and robust. They were late to the market with this feature but have pretty much caught up with the state-of-the-art in one release. With a little more field hardening and a few more features, the Global IP feature may prove to be the a glimpse of the future as regards multi-site load balancing, probably working in concert with DNS Re-Direction, eliminating the site fail-over latency issues of DNS Re-Direction.

4.6.2 Updated Analysis

I humbly admit that of all the vendors, the worst analysis I did was of Foundry in the May, 1999 report. I thought that VRM was an after thought for Foundry. After spending more time with them, it became clear to me that VRM is an integral part of their strategy. Roughly 20% of their revenue is due to VRM-related products. An even more important indicator to me is that towards the end of 1999, Foundry caught up in several areas in which I felt they were lacking – Layer 7 functions, VRM support on their modular platforms, and multi-site VRM. In fact they put so much effort into this, they now do multi-site VRM as good as anybody and better than most.

But Foundry will never be the market leader in this space. The reason is - they don’t need to be. Foundry has much higher revenue than Alteon, for instance, because they have been able to compete and create a strong beachhead in the high-end Layer 2/3-switch market – a market that was roughly 12 times bigger than the VRM market in 1999. So Foundry uses VRM like a bridge player roughing and trouncing the competition. If the customer primarily needs Layer 2/3 switching, position VRM as a differentiator vs. the other offerings in the market. If the customer primarily needs VRM, position Layer 2/3 switching as “tag along” sales to complete the solution. Therefore Foundry gets both more Layer 2/3 wins as well as more revenue per customer for deals that are primarily VRM.

So while Foundry has generally been able to keep up with the VRM market in terms of features, and leads occasionally in some areas, they don’t sit around eating, drinking, and sleeping VRM. They also have to worry about how to compete with Extreme, Cisco, and Juniper in routing. That takes a lot of thinking and a lot of resource. Therefore I expect them to always be a B+ player in VRM, as one of many tactics for remaining an A+ company overall.

4.7 CyberIQ Systems (The Company Formerly Known as Holontech)

4.7.1 New News Since May, 1999

Up until recently, Holontech did not have a lot to announce in the past seven months. Most of their announcements were related to:

- An application effort with Microsoft, where the Holontech technology can be used as the Resource manager (QRM) for an MSCS cluster, allowing the cluster protocols and operation to extend up to sixteen Windows 2000 web servers.
- Low-end 4- and 8-port devices that provide web server load balancing only, and are advertised and plug'n'play in terms of installation simplicity, which I believe includes automated IP address assignment, and automated learning and configuration of web server addresses.
- Co-marketing relationships in the caching (Cache Flow) and VPN (Red Creek) areas.
- The hiring of a new CEO, president, CFO, VP of Marketing, VP of Business development, and VP of International Sales. The announcement of Bill Schroeder as the new CEO was combined with an announcement about the closing of a round of funding (of around \$20M).

Lacking any better information, one (if one is a cynical analyst) would have concluded that Holontech was stalled.

Then, recently, BOOM! Holontech announces:

1. They are no longer Holontech (they are now CyberIQ Systems)
2. As part of their recent financing round, they also announced a technology relationship with Lucent (Lucent Ventures was part of the funding round as well). The nature of the technology is adaptive rules engine capability, which I believe Holontech tends to productize in the form of value-added resource management, looking at events, resources, and performance levels and over time determining “signatures” that identify when load-balancing algorithms should be tuned.
3. An updated mission statement, vision, and high-level product roadmap which, in the short term, catches Holontech up with the rest of the industry in the important areas of Multi-Site (GeoDirect), SSL acceleration, and Layer 7 services and, in the long run, leverages the Bell labs technology to “bring to website resource management in the 21st century what Bell Labs brought to PSTN management in the 20th century.”

4.7.2 Updated Analysis

What a difference experienced Silicon Valley management and an influx of \$20M makes!

Through its birth and even up to now, Holontech has been a very Engineering-driven company. However the new management team will make the company more complete in terms of business strategy, channel development, partnering, marketing, sales, customer support, and all the other non-product related aspects of the success equation.

But even though they have made some great strides technically (in terms of announcements, anyway) Holontech (I'm sorry – CyberIQ Systems) still has a lot of catching up to do in terms of channels and partners. They are doing OK in the Far East, especially Japan, where they get over 50% of their revenue and are one of the leading vendors there. But that percentage is due as much to weakness in the U.S. and no presence at all in Europe as it is to their strength in the Far East.

If you believe my thesis that the year 2000 will be the make-or-break year for vendors in this space, and that the key is not technology but “being there,” and if you are a Holontech fan, then you have to hope for them to ramp up their customer presence extremely quickly. If by 3Q00, F5, Alteon, Foundry and others are having \$25-35M quarters, and Holontech is only at \$10-12M, then the game is over. And even getting to \$10-12 is a huge increase for them. The good news however, is that they are not just starting to invest in the channel. If they were, there would be no hope. They increased their sales staff considerably in 2H99, and started to create partnerships in Europe.

Assuming they can ramp up fast enough to hang in there, Holontech then needs to learn how to leverage their unique capabilities and approaches better. They aren't very differentiated relative to web server and firewall load balancing (no one is). But they have started some unique projects with Oracle and Microsoft. However, they haven't really figured out yet how to leverage those relationships to carve out large niches for themselves, protected both by technology and relationships. If the new management team can do that as well as scale the channel, then Holontech may go into 2001 stronger than they went into 1999 or 2000. If not, they need to look to get acquired.

4.8 HydraWeb

4.8.1 New News Since May, 1999

New Platforms – Hydra 1500 and Hydra 2500

- 2U form factor.
- Available in 1Q00.
- Hydra 1500: \$24,000 for a redundant pair. Six server license. Includes HydraWeb's intelligent agents, cookie features (persistency and preferential services), and URL parsing features.
- Hydra 2500: \$44,000 for a redundant pair. Unlimited server license. Includes HydraWeb 15000 functions, plus multi-site capability, application monitoring, and SSL Session ID tracking.

HydraFire, 11/99

- Integrated application-layer firewall and server load balancer.

Cookie-Based Features, 9/99

- \$10,000 option on the Hydra2000 and Hydra5000.
- Puts a cookie on the client browser depending on the customer group the user is associated with and/or the status of their interaction with the site. Uses that cookie later to regulate access to server resources. As an example of the latter case – paying customers may be issued a higher priority cookie after making their initial purchase.

Application Server Monitoring, 9/99

- \$5,000 option on the Hydra2000 and Hydra5000.
- Agents residing on application servers provide feedback to the scheduler on application health and load.
- *Such server-resident monitoring can be very useful in multi-tier application delivery environments.*

4.8.2 Updated Analysis

From my perspective, HydraWeb is the VRM vendor that gets the least respect. Technologically they are on top of things. They have been first to market with a lot of the innovations of the last couple of years, and when they are not first, they are close.

But in spite of that, I don't know them very well. So I am probably the worst offender at not giving them the respect they deserve. Perhaps that's because they are New York City-based and I've become a West coast guy. But even our field people in the northeast doing end-customer projects don't see or hear about HydraWeb much.

My estimations of HydraWeb revenues and therefore projections of future revenues are essentially just guesses. I could be way off. I need to get together with the HydraWeb people and learn more about them. Until then, take anything I say about them with a grain of salt.

4.9 IBM

4.9.1 New News since May, 1999

Release 2.1 of SecureWay Network Dispatcher

- URL-based scheduling. Requires the use of IBM's WTE proxy server. The proxy server receives requests and directs them to clusters of servers for cgi requests, static content, etc. The Network Dispatcher sits between the WTE proxy server and the web servers and provides load balancing within the clusters. In other words, the proxy proxies to virtual addresses rather than physical addresses.
- Cookie affinity. Similar to the URL-based scheduling described above, requires the use of IBM's WTE proxy server. The proxy server is the entity that resolves the cookie and then directs the request to a particular server.

Release 2.1.1 of SecureWay Network Dispatcher

- Support for running SecureWay on Red Hat Linux Version 6.0

4.9.2 Updated Analysis

IBM's enhancements have been modest. The main enhancements marketed in 2.1 (URL and cookie-based scheduling) are really just application notes associated with the combined use of Network Dispatcher and the IBM proxy server.

I can't really provide much insight into the IBM strategy in this area, because I just don't know. Repeated efforts to find a Product Manager or someone "in the know" to talk to have failed.

I just don't think IBM pays too much attention to this area as a specific product category. Instead, they view it as one of many capabilities they bring to the table to provide a suite of E-Business solutions. In that sense, they may reflect more the future of VRM than the present.

4.10 IPivot/Intel

4.10.1 New News Since May, 1999

Product Renaming (Intel-ification)

- **Intel NetStructure 7110 e-Commerce Accelerator:** Used to be the IPivot Commerce Accelerator 1000
- **Intel NetStructure 7140 Traffic Director:** Used to be the IPivot Intelligent Broker 4000
- **Intel NetStructure 7170 Traffic Director:** Used to be the IPivot Intelligent Broker 7000
- **Intel NetStructure 7180 e-Commerce Director:** Used to be the IPivot Commerce Accelerator 8000.

New Multi-Site Product (September)

The Intel NetStructure 7190 Multi-Site Traffic Director replaces the former IPivot Intelligent Broker 7000M. This isn't just a name change. It's a new product.

- Uses DNS Re-Direction.
- Metrics for "best site" include site load and network latency.
 - Latency is measured by sending out a synchronized message to 7190s at each site – "send a response back (with your V_IP) to the LDNS with this IP address at (for example) 23 hours, 16 minutes, 10.25 seconds." The site whose response gets back to the LDNS first "wins."

New Channel Partners (September)

GE Access, Sun's largest distributor, inked a deal with Intel to represent and support their products. GE Access is one of the few VARs to actually provide some V and some A. They provide eBusiness bundles, installation support, training, and other useful services.

4.10.2 Updated Analysis

The big news from IPivot is that they no longer are IPivot. They are now part of the Network Equipment Division of Intel.

Whenever a big company buys a small fast moving company, I always hold my breath for about eighteen months to see how the integration goes. That much time has not passed yet, so there is still time for Intel to screw up IPivot - but so far so good. In fact, if things continue to go well, this might turn out to be a future business school example of how to execute such acquisitions.

First of all, Ipivot has been pulled essentially untouched into the Network Equipment Division (NED) business unit, newly formed to accommodate Ipivot and focus on Ipivot's mission, strategy and target markets. The GM of this business unit is Brett Helm, the former CEO of Ipivot, and all of Ipivot's R&D, customer support, sales, marketing, and manufacturing operations have been moved intact into this organization. At the same time, carefully selected Intel channels have been introduced to the Ipivot product line and the NED operation has been augmented by some corporate functions like corporate marketing and G&A – things that were just a budget burner at Ipivot.

More interestingly, the Intel NED systems strategy seems to have been heavily influenced by the visions, people, and technology of IPivot. Intel has recently changed their corporate mission statement (which is something they do only about every 15 years) to something like "be the leading supplier of building

blocks for eBusiness.” They wouldn’t have talked about themselves that way just a few short months ago. I’m not sure what is cause and what is effect here, but the bottom line is that IPivot is considered to be critical to Intel’s future success, which is more than can be said for most acquired companies after a few months go by and the honeymoon is over.

So all in all, even though I am generally very skeptical about the ability of large companies to integrate small companies, this one looks like it will work out very well for Intel, IPivot, and their present and future customers. In fact, it might work out so well that if some of the other large networking companies don’t follow suit and acquire & integrate soon, they may find that they gave Intel too much of a head start and they actually end up being a leader in this space!

The Intel Netstructure product line messages are mostly constructed around the issues of Layer 7 services and SSL acceleration. The two are intertwined. The Intel position is that you need to provide SSL acceleration in front of the server farm both to make the servers more efficient and to decrypt the packets so that Layer 7 features can be executed by the VRM function. I tend to agree with this view. Without the need for value-added Layer 7 services in the front-end VRM device, we could have a great a debate about whether acceleration should be implemented per-server (via a card) or as a front-end box. With the requirement for Layer 7, the debate is pretty much won by the SSL Accelerator appliance proponents, such as Intel. But just about every vendor I know of is working on a similar capability to Intel’s, so their time in the sun as the unique supplier of this capability may be short lived, but I’m sure they will move on to bigger and better things as the others scramble to catch up to what they did nine months ago.

Intel also likes to talk about their Intelligent Session Recovery feature, where requests are cached in the VRM device and sent to a different server if the request to the initial server results in an error Return Code of various types. Intel/Ipivot has supported this feature awhile, but so far no other vendor has replicated it. To a certain degree, I believe that this is because IPivot had to spend a lot of the last nine months executing on the Intel acquisition and therefore less of their energy was spent educating the market about this feature than they may have otherwise done. But now, with the Intel marketing might behind them, I expect them to start a full court press on this and the other differentiating features they have in their products and in their development pipeline.

I think you’ll be hearing a lot from Intel in the coming months and years in this space.

4.11 RadWare

4.11.1 New News Since May, 1999

Cache Server Director Enhancements, 9/99

- Support for transparent caches (in addition to the already existing support for proxy caches).
- Industry unique support for making proxy caches look transparent, i.e. running caches in proxy mode, but not requiring configuration of the browsers to use them.
 - This is achieved by having the CSD intercept all traffic and filter it to determine whether it is cacheable or not. The non-cacheable traffic is just forwarded along to the ultimate destination. The cacheable traffic is sent to a selected proxy cache.
 - To send to a selected proxy cache, the Destination Address is changed to that of the destination (NAT), plus the HTTP request is changed from a Relative URI (as generated by the source browsers, who don't know they are accessing a cache) to an Absolute URI (as required by the proxy caches) by looking at the host header later in the HTTP request. This is the tricky part.
 - Why is this useful? Some caches can't run in transparent mode and some that can run slower. This feature of the CSD allows ISPs (who require transparent mode operation) to have a wider selection of choices for their caching solution.
- Support for re-direction of multiple (up to 20) configured TCP Destination Ports
- Support for multiple policies to be configured per IP address or destination URL. Policies include:
 - Blocked: Access to the specific URL or IP address is blocked.
 - Direct: The IP Address or URL is not cached. Requests are sent along directly to the destination.
 - Local Server: The request to the IP Address or URL is sent to a local cache.
- Full NAT.

Firewall Load Balancing Enhancements, 9/99

- Enhancements to the RadWare Fireproof product, which is a product with features specifically suited to Firewall, VPN device, and router load balancing. (Other vendors have these features – of some subset of them – as part of their core product. I think it's a pretty good idea to split them out in a product people can buy and configure specifically for these applications).
- Integrated support for both transparent and proxy firewalls. (Previously, proxy firewall support was available on the WSD, but not Fireproof. This is a good addition to Fireproof capability so that one box can be used to load balance a cluster of firewalls that may have both transparent and proxy services turned on).
- Industry unique Port Rules feature, allows one to configure the box as multiple logically separate boxes, to avoid the "firewall sandwich" box proliferation problem.
- Support for firewall "groups," with each group responsible for different applications or source/destination address ranges.
- Load balancing by packet and byte rates.
- Outbound NAT (IP Translation).

Multi-Link Load Balancing, 10/99

- Oriented towards balancing loads and ensuring high reliability for sites served by multiple WAN links from multiple ISPs.

- Eliminates the need for complex routing protocols or coordination between the ISPs to use their WAN links (and networks) effectively.
- Latency, router hop, and load tests are performed regularly over each ISP WAN link.
- The latency test can be configured to test various downstream points (e.g. ping to the aggregation router at the POP, ping to a server in the ISP cloud, ping to the BGP edge router) so that the source of delays can be deduced.
- Outbound, the source IP is modified to be consistent with the IP addressing each ISP is expecting.

I like this capability a lot. It is elegant, yet straightforward. The LinkProof box has many of the features of FireProof as well, so you don't have box proliferation when performing firewall and link load balancing at the same site.

Partnerships

- CheckPoint (OPSEC Certification)
- WebTrends
- Lucent
- Network Engines
- Network Associates (FLB)
- AXENT (FLB)
- Lucent (FLB and part of their VPN alliance)
- Cylink (FLB)
- Secure Computing (FLB)
- Starburst (content distribution)

4.11.2 Updated Analysis

RadWare is typical of why I like spending time in the VRM space. The key marketing and technical people are excellent to work with and it's fun to watch them identify and solve interesting problems.

RadWare is another company that had a successful IPO this past fall. I was happy to see it go well because I like the RadWare people. Oh, and also because they included me in their Friends and Family program.

RadWare does make me angry in one important sense, however. They are too easy to work with. When they are about to come out with a new product, they call me before hand and tell me "this is what it was designed to do, this is how it works, this is what it wasn't designed to do and therefore doesn't do well yet. Oh, and by the way, we call it the widget load balancer, because it load balances widgets." This is all usually accomplished in a 30-minute phone call. But I feel empty and unsatisfied when the call is over. Where are the over-arching mission statements, the outlandish claims, and the unsubstantiated competitive knock-offs, all wrapped up in newly minted esoteric terminology? Don't they know that half my fun comes from cutting through all that crap! Sigh...

As much as I like these people, I'd like to see them grow their company a bit faster. Although they are profitable and continually growing, I fear they are losing ground to some of their highflying competitors. Quarter-by-quarter growths in revenue of \$1-2M seem paltry when you have competition that is growing by \$4-6M per quarter. And even more so when you consider that RadWare is one of the veterans in this space. They have had plenty of time to find their legs. In 1998, F5 Networks and RadWare had almost exactly the same revenue. Not nearly true in 1999.

What's the problem? The products are excellent, their margins are fantastic, and they know their stuff.

One problem is that people (at least this person) has a difficult time “labeling” them such that they carry with them a simple and unique positioning. They are not “the Content Switch Company,” or “the Web Data Center switch” company, or even “the distributed software company” or the “ITM market leader.”

To me, they are the “Israeli-based industry veteran with good products in all ITM product categories whose products are solid and low priced and sold globally through a wide variety of channels” company.

Can RadWare find their legs and establish a stronger position in this marketplace? I think so. I’d be less optimistic if they had technology or margin problems they had to solve. And I know they can execute because they have focused on and developed the whole Firewall Load Balancing concept almost single handedly. Investments they made in Sales and Marketing, and in Firewall Load Balancing market conditioning in 2H99 should start to pay off in 2000.

But – I think before they have a chance to give it the good fight too much longer, someone with channel but without product will see them as a good deal and swoop them up. Their technology, product margins, and people, are just too enticing to pass up. For their sake, I hope if that happens they find a good home.

4.12 Resonate

4.12.1 New News since May, 1999

Local Load Balancing

Central Dispatch 3.0, 12/99

- Cookie-based scheduling for persistence and preferential services
- Threshold-based scheduling

Dispatch-SLB NOW™

- Free SLB!
- Download from their web site.
- Support for 2 servers only.
- Basic Layer 4 SLB plus URL-based scheduling. But no cookies-based persistence or preferential services, no threshold policies, and no SSL Session ID persistence.

Resonate Commander 3.1 (March 2000)

- Integrates external probing and site query functions into the core product. These tests examine the operation of every tier of complex application delivery architecture. Previously, many of these functions had been achieved through user integration of resonate with products from Freshwater.
 - TCP Port Test.
 - URL Transaction Test; allows session information to be verified. Useful for sessions that must persist across multiple pages.
 - Multi-URL test; Useful for ensuring referential integrity among pages.
 - HTTPS Test.
- Tools for developing customized tests.
- Support for more complex actions.
 - Expands current URL string matching by supporting logical operators such as AND, OR, and NOT.
 - Linked actions for when an event occurs and a “return to normalcy” event when the condition goes away.
- Threshold configured rules.
 - Invoke actions based on min/max rules applied to any Central Dispatch statistic.

Multi-Site Load Balancing

Global Dispatch 3.0

- Java-based GUI for configuration and monitoring.
- Weighted load balancing.
- Enhanced static scheduling – support for tables up to 100,000 entries large.
- DNS Reverse Look-Up – site selection based on parsing the domain name of the requesting LDNS.
- Persistent table synchronization for redundant units.

Expanded Environment Support

- Linux.

- IBM Netfinity 4000R.
- Thin IBM RS/6000 B50.

Managed Services

- Outsourced remote management of web sites from the Resonate Service Operations Center (SOC).
 - InterAct Services: 24x7 monitoring and control service option essentially outsources all day-to-day operations.
 - ProAct Services: Provides monitoring and management oriented towards maintaining service levels via capacity planning performance tuning, establishing preferential service policies, performance monitoring, and on-going tuning.

Partnerships

- Siebel (OEM)
- eBusiness Technologies
- Vignette
- Whale Communications
- Penguin Computing
- Sun Professional Services
- Sterling software (OEM)

4.12.2 Updated Analysis

Viva la difference. Resonate's biggest possible strength -- and possibly biggest historical problem -- is being different. They have a different and unique architectural approach than any vendor, they have a different sense of target customer, and they have a different business model.

One of the early pioneers of VRM, Resonate's approach is a distributed software model, where two software components reside on each web server in a cluster. One software component is an agent that oversees the health and load of the server, the other is a portion of the scheduler function which is installed in the protocol stack just below the IP layer. The second component is why I call the Resonate approach distributed software, because the scheduler function is distributed between (essentially) in-bound tasks and outbound tasks. Since the outbound tasks are handled on each individual web server, the in-bound scheduler is off-loaded and therefore has increased capacity. It is kind of like the Direct-Path-Return option supported by most of the hardware vendors, with the important difference that Resonate can perform Delayed Binding functions in this mode because they can "touch" the packet on the return path without it having to pass back through the scheduler. No hardware vendor can do this. And in fact, Resonate was the pioneer of URL-based scheduling, supporting that function for a good 1½ years before any other vendor did.

There are pros and cons to the Resonate approach, however. The good news is that it eliminates the need to introduce new boxes into your environment, can overlay any physical and logical topology you have established for your site, and scales pretty well, which is especially noticeable if you turn on advanced Delayed Binding features. The bad news is that you have to install and maintain the software on all your servers, and if the software isn't supported on your OS or network card, then you can't use the Resonate solution.

The hardware vendors have been hammering on those negative attributes for some time now, and they generally win the argument if they are talking to someone with a networking background who views the VRM problem as a networking problem. Who wants to deal with software on the servers if you stick a couple of boxes in the packet path and obtain the same capability? And historically, Resonate hasn't helped themselves out in this regard. The products were difficult to install, they didn't support a wide

range of environments, and their agents, which ate up some CPU cycles, didn't do much in terms of adding value to the solution.

As a result, in the SLB wars of 1998, Resonate didn't fare all that well. During that period networking people were making most of the vendor selection decisions and most people didn't appreciate the power of Delayed Binding and more specifically URL-based scheduling. So Resonate's unique advantages weren't helping the much. They were struggling, although they did have some pretty impressive marquee customers.

This all changed in 1999. New management came in and invigorated the company; aggressive hiring resulted in critical mass in engineering, customer service, sales, and marketing. Most importantly, Resonate took a step back and re-evaluated their strategy. Basically they asked themselves, "where are we succeeding and where are we not?" The answer was that their sweet spot is not in the small cluster deals, but in large-scale applications, especially where the decision maker comes from the systems/applications world, rather than the networking world. That kind of decision maker, who is used to instrumenting applications and back-ends with systems and resource management solutions like those available from Tivoli, CA, or BMC, look for solutions for internet traffic management which have a similar look and feel. So they actually embrace the distributed software approach and eschew the "foreign box" network-oriented approach.

So Resonate has re-oriented their company direction towards those decision makers and those applications. As systems management vendors have discovered in the past, the requirement here is not necessarily a "one size fits all" shrink-wrapped product, but a set of tools with enabling functions that can be flexibly integrated into complex systems to add value. That implies a customer support model that is as much about customization services and outsourcing as it is about product. Consistent with that, Resonate invested heavily in Professional Services and Remote Management Services in 1999.

As Resonate has shifted their focus to high-end deals, they have reduced their focus on small deals. In fact, revenue from the basic Server Load Balancing function is no longer important to them, as evidenced by the fact that they have recently made available a free version of such which is downloadable from their web site. What they want to do is seed the world with their basic technology, and then come back to people with up-sell propositions related to more advanced products, professional services, and outsourced services.

The shift of strategy seems to be working. Resonate has submitted their S1 as they head towards IPO. In examining that, you'll see increasing momentum throughout the year, but especially in the 2nd half of 1999. That uplift in business is largely due to a focus on some very big high-end deals – deals that result in pretty good initial revenue but are even more interesting because they imply a long-term relationship with the customer related to new products and a variety of on-going services. Most of the hardware VRM vendors say "we never see Resonate." And I believe that is generally true. They just don't get invited to the same parties. While the hardware VRM people are selling to site architects about cluster solutions, Resonate is talking to CIOs and business managers about complete enterprise and system solutions. It's a different target application/customer set. So different I am almost ready to declare that Resonate is not a VRM company any more and should be lumped into the Systems Management category. But since VRM is at the core of their added value, I won't do that yet.

So what is going to happen with Resonate? Although they started to turn around their business in 1999, they are still not a dominant player yet. And although Resonate thinks their model will be the dominant VRM/systems management model of the future, I don't think that's true. There are already too many other approaches out there and running well for the market to shift significantly to the distributed software approach. Instead, I think Resonate will establish a strong and defensible niche at the high end. Success

or failure will depend on whether they can expand that high end to address 25% of the applications or whether it contracts to 1%. The wild card in that respect will be the way in which enterprises adopt VRM technology in the coming 12-24 months. Much of the market growth for the next few years will be driven by enterprise adoption of VRM. If the decision makers in the enterprises are the people who were responsible for the deployment of Tivoli in the past, Resonate is in great shape. If the decision makers are the people previously responsible for deploying the Cisco routed network, they are not.

Personally, I think an increasing percentage of people will address the problem as a systems management problem than a networking problem, which bodes well for the future of Resonate because if they do, it eliminates almost all the other competition discussed in this document, with the possible exceptions of IBM and F5 (who I view as essentially a software company with an expanding systems view similar to Resonate's in many respects). What could slow Resonate down is lack of clarity about where they fit in the world. Are they Server Load Balancing? Are they Service Level Management? Are they Systems Management? But – confusion is where the big market opportunities lie. We'll see how it goes.

I expect their IPO to be quite successful. I'm sure hoping so - Acuitive owns some shares from work I did with them in 1999.

4.13 Others

4.13.1 The “F5 Enabled” Layer 4 Switches

F5 Networks has engaged in OEM/reseller and technology licensing agreements with Cabletron, 3COM, and Extreme. The end result is that all of these companies have imported F5 Networks Layer 4 SLB capability into selected platforms and are also selling the rest of the F5 Networks product line as needed into their customer base.

The technical and business details vary a bit from vendor to vendor.

3COM

In November of last year, 3COM and F5 Networks announced an agreement to port F5 Networks code onto an Intel –based blade for the CoreBuilder 9000. 3COM would also resell the rest of the F5 Networks product line.

Recently, 3COM announced plans to discontinue sales of the CoreBuilder 9000 platform and to totally get out of the enterprise market.

Whoops!

Cabletron

The Cabletron announcement was the first of the announcements F5 Networks made with switch vendors, going back to May 1999. This announcement was purely an OEM deal, where F5’s Big/IP product would be re-labeled by Cabletron and sold as a bundled solution with their Layer2/3 switches.

Since then, Cabletron has announced support for some Layer 4 VRM functions native to their SmartSwitch Product Line (there is one version with 24 port 10/100 and 2 GigE ports and another with 8 GigE ports). The feature set is modest, but the price/performance is excellent.

Extreme

Probably F5 Networks and Extreme Networks made the most important of the switch vendor announcements in August 1999. The companies announced the intention to port Layer 4 F5 Networks code into the Extreme switches, and for Extreme to resell the rest of the F5 Networks product lines as enhanced solutions around that foundation. In fact, Extreme has already started shipping the Summit 71 product, which integrates some of the F5 Layer 4 load balancing functions. This gives F5 Networks a great expansion of their channel and gives Extreme the ability to come closer to parity with Foundry and Alteon in regards to VRM features on Layer2/3 switch platforms.

4.13.2 Lightspeed Systems

Lightspeed Systems is a company based in Bakersfield, CA, focused on software solutions that run on Windows NT/2000 environments to provide ITM functions at three points of the system:

1. The gateway
2. The server
3. The desktop

Their products are:

- **Traffic Control for e-Business™** : Provides a “hardened” IP stack, firewalling, NAT, SLB, and traffic statistics monitoring (\$9,995).
- **QoS Control for e-Business™** : Adds policy-based traffic prioritization, via classification and tagging functions (\$3,495).
- **Server Control for e-Business™**: Monitors services (such as SQL, FTP, NNTP, etc.), sends notifications when problems occur, and initiates service re-start or system re-boot when appropriate (\$995).
- **Desktop Control™**: Personal firewalls and a framework for ISVs to create Internet plug-ins (free).

The main value-add of Lightspeed is probably not the functions and precuts above in and of themselves, but tight integration into the Windows NT/2000 system environment. Auto-learn of pre-configured Windows users and groups, configuration wizards, drag-and-drop procedures, and other UI aspects familiar to NT users are leveraged throughout the Lightspeed product line.

4.13.3 Phobos

Phobos is a Salt Lake City-based company that has been around for a while, mostly providing engineering services and Intellectual Property on a contract basis for customers such as AMD and 3M. Last fall, they introduced a PCI-card version of a Layer 4 Server Load Balancer (IN-Switch). They have since followed that up with a PCI-card SSL Accelerator (IN-Boost) and appliance versions of both products (IPXpress and SSLXpress).

The Phobos products are simple to install, low cost, and reasonably featured (for Layer 4 devices).

I like the fact that Phobos provides SSL acceleration options both in the form of PCI cards and front-end appliances. One can achieve different levels of scalability depending on the need. The appliance also enables Layer 7 features on VRM devices to be used (but not their VRM devices, as they are Layer 4 only).

I don't care too much for the VRM PCI card – generally I don't like to open up a server to provision a networking device - although I could see an application in a rack of equipment that is packaged for colo deployment, to save rack space.

4.13.4 StoneSoft

StoneSoft is a Finnish company that provides software solutions for eCommerce. The company has about 350 employees.

In 1999, StoneSoft introduced software solutions under the name StoneBeat FullCluster that co-reside with CheckPoint, Network Associates, and AXENT firewalls, providing High Availability and Firewall Load Balancing.

In February of this year, StoneSoft introduced expansions to the StoneBeat product line that provide web server load balancing and cache server load balancing.

Sun Professional Services is tightly aligned with StoneBeat and deploys their firewall load balancing solution in many of the architectures they deploy.

Products Summary

Their firewall load balancing solutions support the three major vendors of software-based firewalls – CheckPoint, Network Associates, and AXENT. They have no solution for appliance-based firewalls such as Cisco PIX or Netscreen.

- FullCluster for Firewall-1
- FullCluster for Gauntlet
- FullCluster for Raptor
- WebCluster
- CacheCluster

The StoneSoft product line is pretty mature – supported on a large number of hardware platforms, O/Ss and NICs.

Architecture

In all cases, the product architecture is similar. It operates a lot like the Envoy product that Microsoft bought from Valence as source to create their Windows Load Balancing Service (WLBS) feature. By that, I mean it works using Layer 2 Multicast so that all incoming and outgoing packets are sent to each firewall (or server). Each device supports the same IP address and MAC (multicast) address. Devices surrounding the firewalls, such as routers and servers, are provided the Layer 2 Multicast address as a response when they ARP for the cluster's virtual IP address.

The NIC drivers on each firewall (or server) then filter each packet to determine which packets should be dropped (because they are being serviced by some other device) and which need to be passed to higher layers of software. I believe that a hashing is performed on IP SA/DA pairs (possibly TCP ports as well) and the possible hash results are distributed between all of the firewalls (servers) so that the load is statistically spread between the devices. If a device fails or become over-loaded, some or all of its assigned hash results can be re-assigned to another device. Note that persistency is not maintained in this process.

I believe that the hash results will be the same in both directions of traffic flow, to ensure that the same firewall is processing the traffic in both directions. This is important in many stateful firewall situations.

The StoneBeat products have good techniques for testing firewalls, NICs, and downstream devices to identify firewall and path failures.

One nice thing about being a software solution is that you don't have the "firewall sandwich" issue that you have with hardware-based FLBs. This makes the solution topology *much* simpler than in the hardware case.

The StoneBeat approach is simple and lends itself to reasonably good load balancing and very good fail over characteristics. The downsides are that every packet needs to be forwarded to every device in the cluster. That means that LAN switching cannot be used to increase the effective throughput of the cluster. If each device is connected to a Fast Ethernet connection, the maximum cluster throughput is 100 Mbps even if each device is connected to a dedicated switch port. Also, the NIC driver must spend time processing each packet, even those that are not to be processed by that particular device, which undoubtedly degrades perform to some extent.

Another downside is that the solution can only be used to load balance software-based firewalls. You can't use it with appliances such as Cisco PIX, Nokia, and Netscreen.

Finally, the feature set is Layer 4 only.

Architecture Summary

- Layer 2 multicast.
- Each firewall NIC looks at every packet.
- They communicate with one another to determine who will handle which flows.
- This has the advantages of:
 - The firewall cluster is represented by a single IP address and a single MAC.
 - Bi-directional traffic can be forced through the same firewall.
 - It is easy to define “roles” for each firewall and then change them without having to modify any external components.
- This has the disadvantages of:
 - Doesn’t scale well. Every firewall NIC sees every packet.
 - Layer 2 switching doesn’t increase LAN bandwidth – multicasts have to be broadcast throughout the VLAN.

For the most part, I would view StoneSoft’s StoneBeat solutions as a low cost, easy to deploy High Availability solution for a small cluster of devices, especially for software-based firewalls.

4.13.5 Rainfinity

Rainfinity is a start-up company in Mountain View, CA. Their main claim to fame is clustering technology sourced from Caltech/NASA. Similarly to StoneSoft, Rainfinity provides a software-based solution for firewall load balancing.

Products Summary

Rainfinity presently only has one product – firewall load-balancing support for CheckPoint Firewall-1 and VPN-1 products. If you are only going to support one vendor, that’s the one to focus on.

Architecture Summary

- Firewalls externally represented by a “pool” of Virtual IP addresses.
- Each firewall owns one or more V_IPs.
- Routers and servers are configured with the V_IPs in their list of default gateway addresses.
- Load balancing is achieved by having different routers and servers use different default gateway addresses (V_IPs)
- When a firewall gets overloaded, one or more V_IPs on that firewall are moved to a less busy firewall.
- The new firewall that now owns the V_IP issues a Gratuitous ARP, so that external systems will learn the new MAC address for the V_I and direct traffic to it.
- This has the advantages of:
 - LAN switching can be used to increase LAN bandwidth because the large majority of traffic is unicast.
 - Firewalls only have to process the traffic directly sent to it.
- This has the disadvantages of:
 - Bi-directional traffic cannot be passed through the same firewall. The firewalls themselves must support state synchronization between different firewalls (this is one reason why I believe Rainfinity only supports CheckPoint), but turning on that feature can be very resource intensive on a firewall.

- Load balancing granularity and control is poor.
 - It is left up to luck whether the external devices spread their default router selection across a large number of V_IPs or use a concentrated few.
 - You can't balance individual flows or individual users. When you move a V_IP, a potentially huge block of traffic moves with it.
- The administrative burden of managing a large number of V_IPs is high.

4.13.6 WebManage

WebManage is a New Hampshire-based company that has been in existence for about three years. Their initial focus was on management tools based on web server log analysis, a product they continue to support and enhance. They have also gotten into the VRM business via their product called i-Scaler, for which they announced Release 2.0 in March, 2000.

i-Scaler is a software-based VRM solution. Similar to the IBM solution, you buy the software and load the scheduler portion onto an NT platform to create an appliance yourself. Thus you get to choose price/performance. Use the latest and greatest Intel platform or save some money and use last year's. Then (optionally) you can install lightweight agents onto the servers being balanced. These agents do some local health checking and also feedback information of load parameters such as CPU utilization. Agents are supported for NT, Solaris, and Linux.

i-Scaler is a low cost solution, costing only \$4,995 for a license for up to five load balanced servers. (You need to add in the cost of the NT platform to get a complete cost picture).

Feature –and-function-wise, I-Scaler performs MAC Address Translation and is a Layer 4 device. WebManage has emphasized some Class-of-Service features (CoS) in release 2.0, which allow you to allocate server resources dynamically to different hosted customers (identified by V_IP), transacting customers vs. browsing ones (identified by HTTPS vs. HTTP) or by user (identified by Source IP Address and/or Source Port).

I don't think the present CoS functions supported by WebManage are all that useful. To the modest degree that I think CoS features are important (which is that I like them as background threshold based mechanisms at sites where the complexity is really worth the effort), I think that the box needs to become content-aware (Layer 7) to really provide the granularity of user and traffic classification required to leverage CoS techniques.

Until that happens, I think that WebManage should be viewed as a low cost, easy to manage solution for basic VRM applications, similar to CoyotePoint, Phobos, the Layer 4 switches, and most of the larger vendor's low-end solutions.

5 Some Background Information

5.1 Who Is Acuitive?

Started in 1997, Acuitive is a strategic consulting firm focused on the development and application of emerging networked computing technologies. Acuitive provides a wide variety of technical and marketing advisory services to equipment vendors, service providers, and enterprise network planners. Acuitive also publishes qualitative research reports on emerging technologies, to educate end-users and stimulate market development.

The Acuitive Network Planning and Management Group assists enterprise network clients with strategic planning for IP services. The NPMG services address two key needs in enterprise network planning: Network Requirements and Application Requirements. The services include baselines for networks and applications, certification for new application rollouts, capacity planning, design of complete network management systems, advice on Virtual Resource Management and WAN bandwidth management strategies.

The Acuitive Business Strategy Group helps vendor clients with market strategy, business cases, target customer requirements, product line plans, product and service requirements, alliance partnerships, competitive evaluations, product and program management, company and product positioning, and outbound marketing.

We've been involved in the development and application of VRM technology from day one. As technologists, VRM intrigues us because it operates at the convergence point between applications and networks. As advisors to end users, it's critical because it is an important technical underpinning for making new businesses and business processes based on web technology more reliable and scaleable. As strategists consulting to vendors, it's a technology that has spawned new product categories and companies, and will be a key component of the IP services functions larger vendors need to offer to rise out of the increasingly commoditized network "plumbing" market.

As a result, we've been involved with or have observed hundreds of deployments of the technology and we have hands-on experience with almost all the products available on the market. We have also consulted with several of the vendors to help them orient their products to emerging customer needs.

5.2 About Our Reports and Subscription Service

We've "bottled" our experiences in this area to help you get educated on the technology issues of VRM, characterize your application, and choose the right vendor solution for your application. To that end, we have published two comprehensive research reports.

One, entitled *Virtual Resource Management: Key Technologies, Tricks of the Trade, and Application Requirements* provides a detailed tutorial of the various approaches to building VRM solutions. The values of various techniques and features available in the market, as related to specific application types, are discussed. The result is a "hit list" of key attributes to look for in a VRM solution for your particular application. The attributes are organized by: **policies, mechanisms, feedback, performance, redundancy, and management**. These areas of consideration are the key aspects of a VRM system to evaluate when architecting a solution. Some areas of technology, such as preferential services and security, are still changing rapidly. So those who acquire this research report will automatically be enlisted into a subscription program providing updates (such as this document) and papers on a technology area or application note of interest through April 2000.

Next you need to choose a specific vendor to implement your solution. To help in these considerations, we also offer a second research report entitled “Virtual Resource Management: Which Vendor Is Right For You?” This report summarizes the capabilities available from each of the key vendors in the market today and maps their capabilities against application requirements to create a “short list” of vendors and products to consider for various types of applications. This information is always changing, so those who purchase this report will automatically be enlisted into a subscription program providing quarterly updates (via e-mail) through April 2000.

For more information on the research reports, go to www.acuitive.com. There you will find additional detailed information, including a complete Table of Contents, for each report. These reports are orderable from the web site.

Please direct any comments, questions, opinions regarding VRM to vrn@acuitive.com.

5.3 About The Author

Mark Hoover is the President and co-founder of Acuitive. Mark worked at AT&T Bell Laboratories for about ten years, and was involved in the development of satellite transmission and fiber optic devices and systems, high speed packet switches, and LAN products based on emerging 10-BASET and FDDI standards. Mark also ran the team that provided technical support for the AT&T OEM agreement with Cisco. In 1990, Mark left AT&T to join SynOptics. At SynOptics, Mark was initially responsible for the definition and development of the (at the time) new generation hub platform – the System 5000 (although Jim Vogt did most of the hard work). Mark went on to form the internetworking product line at SynOptics, which was influential in the merger with Wellfleet to create Bay Networks. At Bay Networks, Mark formed the Internet/Telco Business Unit to define, develop, and market products for the service provider community.

Mark formed Acuitive at the beginning of 1997, along with Dave Danielson. Mark now spends his time running Acuitive (which doesn't take much effort), studying technology and market trends, providing strategic consulting advice to vendors in the general area sometimes called “IP Services,” and acting as a “trusted advisor” to several companies at the CIO level. Mark stays in contact with almost all of the key vendors in the VRM space and provides consulting advice, both formally and informally, and learns a lot in return.

5.4 Thanks From The Author

As the readers of this and other VRM documents from Acuitive may know; following and commenting on the VRM industry is a labor of love, performed during times when I can't find real work. In that light, most of my thanks go to the representatives of the various vendors I interact with who patiently teach me and kindly suffer my errors and misconceptions. In response, I pledge fairness, meaning I endeavor to make mistakes about *every* vendor. Some day I'll get it right.

I would also like to thank the employees of Acuitive who work with the products discussed in this document and who actually wrote this document. Mostly I just dry-walled some words onto the framework they built. Related to this, Dave Logan deserves special credit for working with me late into many evenings even while the beer was getting warm.