The Leader in Application Traffic Management

F5 Networks Ltd
EMEA Headquarters
Clarke House
65 High Street
Egham
Surrey TW20 9EY
UK
Tel: +44 (0) 800 587 2233
Fax: + 44 (0) 1784 497 211

F5 Networks GmbH
Einsteinring 35
85609 München-Dornach
Germany
Tel: +49 (0) 89 94383 0
Fax +49 (0) 89 94383 111

F5 Networks SARL
Le Sésame
8 rue Germain Soufflot
78190 Montigny-le-Bretonneux
France
Tel: + 33 (0) 1 39 30 38 90
Fax: + 33 (0) 1 39 30 38 91

F5 Networks
Barcelona Business Center
Passeig de Gracia, 21, 1-1
08007 Barcelona, España
Tel: + 34 93 467 1309
Fax: +34 93 467 1310

F5 Networks
Hardwareweg 4
3821 BM Amersfoort
Postbus 1466
3800 BL Amersfoort
The Netherlands
Tel: +31 (0) 3345 466 11
Fax: +31 (0) 3345 466 66

F5 Networks
Liljekungsvägen 184
Hässelby S-165 75
Sweden
Tel: + 46 (0) 7057 820 21

www.f5.com

CONTROL YOUR WORLD

October 2003

Application Traffic Management:

Your Guide to

Delivering secure, predictable and cost-effective applications

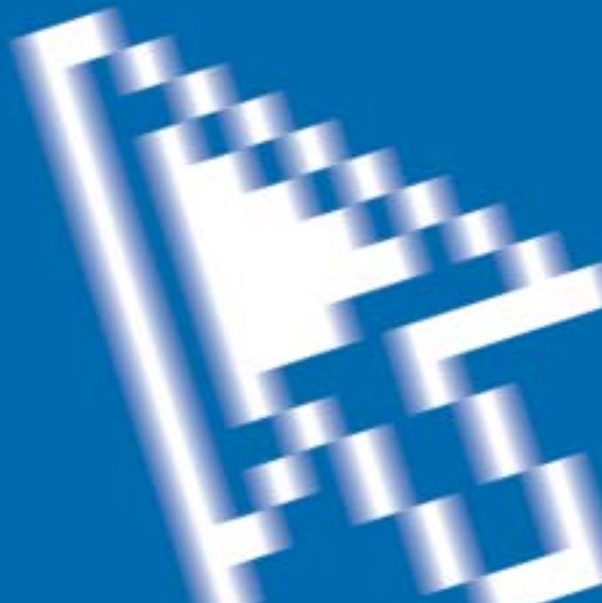

Supporting Online Business

Information Technology (IT) infrastructures have become increasingly dispersed to meet the business needs of geographically diverse end-users and their trading partners. As a consequence, the requirement for structured integration between business application software and internet-based networks has never been greater. This integration is critical in order to secure expected business efficiencies from IT investments and provide the levels of service demanded by employees at work, at home or on the move.

This guide describes how recent innovations for the handling of application data across computer networks can lead to significant improvements in the operational efficiency of organisations.

Application-Aware Networks

The IT industry has responded to the changing working environment of its customers with a number of technology initiatives.

Enterprise applications from nearly all the Independent Software Vendors (ISVs) like Microsoft, Oracle, SAP, BEA, Siebel, Peoplesoft, IBM, Sun and Webmethods are currently being 'web-enabled' as they become more network-centric.

Web Services – defined as the ability to integrate different systems, applications and organisations regardless of the hardware, operating system, application type or location using industry standard protocols – are being rapidly developed to offer integrated online software services both inside and between organisations.

Mobile Computing – meaning the ability to deliver any application to any type of mobile device at any moment in time – is becoming a core business requirement to support the needs of an increasingly mobile work pattern and meet individual lifestyle goals.

The challenge to being able to deliver these applications and services successfully lies in providing high system availability, application-level security and simplified implementation and maintenance in order to achieve operational efficiencies.

In the past it has been possible to achieve these goals for internet-based HTTP, web-traffic only. Now, with the advent of **Application Traffic Management** pioneered by F5, it is possible to extend these capabilities to all Internet Protocol (IP) based applications and services, enabling significant business benefits. And crucially, this can be achieved relatively simply without the need to re-write costly application software code or replace legacy computing resources. The key is to be able to offload certain communications and security functions from business application software whilst being able to inspect all the content flowing over the network between applications and users.

Such 'content awareness' coupled with the ability to automatically apply business rules and policies to network traffic means that intelligence can be applied to the handling and management of IP traffic. This in turn can lead to improved use of current and planned IT resources, with less error-prone manual intervention and a more business-oriented, application-aware network.

Developments in Traffic Management

To appreciate the significance of these innovative developments, it is worth exploring the evolution of traffic management techniques over recent years.

In the mid 1990s, as the internet explosion unfolded and user demand for the serving of web pages grew, businesses faced the challenges of scaling their web sites. Initially, the deployment of bigger, more expensive servers was seen as expedient but it was quickly realised this was a costly, inefficient solution that exposed the organisation to a single point of failure. Implementing smaller groups of less expensive servers running a common HTTP or web-based application became a preferred solution. In this scenario, traffic could be dynamically distributed across the group of servers while making them appear as one server to the user making the request. Thus the concept of Local Load Balancing was born, representing the first iteration of traffic management. Load balancing allows IP traffic to be more efficiently distributed, offers greater economies of scale and provides greater fault tolerance.

In the late 90s, businesses realised that they needed more than simple local load balancing. They also needed to monitor the health of servers, content and applications. This allowed them to make intelligent decisions on where best to route traffic to optimise performance and content availability. The development of Application Health Monitoring ensured that users were connected to the freshest content served from the most responsive part of a web site.

Also around this time, web sites started to be distributed geographically to help improve regional response times, overall site performance and enhance fault tolerance capabilities. In many cases caching technology was employed as a cost-effective method of deploying content closer to users to enhance response times, improve content availability and help reduce bandwidth costs when distributing or accessing fresh content. In order to effectively coordinate such content located across geographically distributed sites, load balancing techniques were extended to create Global Load Balancing capabilities.

Balancing traffic evenly between low cost servers was a primary requirement as web sites proliferated, but as many businesses started to move more business critical applications online and e-commerce became a reality, the need for security in the form of Secure Socket Layer (SSL) Encryption became a critical requirement. To prevent the complex mathematics of encryption algorithms slowing down server performance, SSL technology was built into Internet Traffic Management (ITM) hardware by the more enlightened vendors to improve performance.

Internet Traffic Management has continued to grow in popularity with businesses because it helps achieve economies of scale and has proven to be the most effective architectural design for high availability, scale and security. It has grown beyond a means to route traffic to an available web server and now represents a way to intercept, inspect, transform and direct traffic for core business applications to the correct resource, based on specific business policies.

The Role of Application Traffic Management

Enterprises and Service Providers now wish to extend the benefits of traffic management to all of their IP-based applications. They want to extend the flexibility of the current best practice architectures to all back-office applications in order to achieve the same operational efficiencies they have seen with web traffic. The intelligence and flexibility needed to support these processes is enormous.

New Drivers of IP Traffic

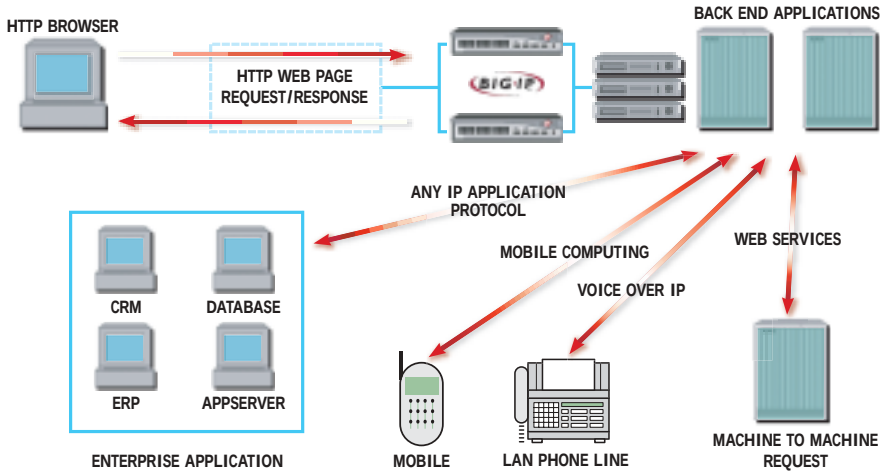

A new way of thinking had to be developed alongside a new architecture in order to support the different types of applications and services traversing the network such as streaming media, web services, video, voice, enterprise and mobile applications.

Application Traffic Management is this new approach. The challenges that need to be addressed by this approach include:

- Ensuring Quality of Service and Manageability
- Applying business policies and rules to content delivery
- Supporting increasing traffic volumes
- Remaining flexible and future-proof
- Delivering Enterprise applications securely
- Delivering operational efficiency and cost control

BIG-IP software from F5 is the first Application Traffic Management capability on the market that can process any application or Web Service, ensuring quick response times, reliable sessions, easy scalability and application-level security – all in a single solution.

How does Application Traffic Management Work?

The critical functions performed by BIG-IP to enable Application Traffic Management techniques to be employed can be summarised as follows:

Deep content inspection of application data traffic traversing the network

Traditional traffic management techniques allow only the network data packet headers, or address information, to be read. Application Traffic Management techniques allow both the address header and the content payload to be inspected. This enables more intelligent handling of the application data. Unique technology developed by F5 called the 'Universal Inspection Engine' makes this possible.

The ability to apply business rules and policies to the handling of network-transported content

The handling of application data can be determined by automatically applying pre-determined instructions to the Universal Inspection Engine. These instructions take the form of programmable scripts called 'iRules' that reflect specific business policies or rules for handling specific types of application data.

Effective integration between business applications and network functions through easy-to-use software tools

F5's BIG-IP communicates with other traffic management devices and application software through a standards-based interface called 'iControl'. This interface is based on the same standards as rapidly emerging Web Services and allows an integrated solution to be created in a timely and cost-effective way.

Security enforcement

The ability to inspect network traffic content means that security threats can be identified and a centralised application security policy enforced. Thus specialised Intrusion Detection Systems can be enhanced to become Intrusion Prevention Systems. F5 calls this technique 'Dynamic Security Control'. BIG-IP is also able to offload Secure Socket Layer (SSL) security functions from the servers for enhanced security and performance.

System performance that meets operational business requirements

Deep content inspection, SSL encryption/decryption and the application of iRules require significant processing power. F5 has pioneered a customised and highly advanced hardware and software architecture to meet the demands of Application Traffic Management.

What do the Analysts say?

Most analysts agree that it is critical to make the creation and delivery of any IP-based application or service or web service predictable, secure and cost-effective regardless of the network environment. Analyst Gartner shows the progression of Content Networking in a hierarchy of needs, displaying the steps organisations take to deliver content.

Hierarchy of Needs for Content Networking

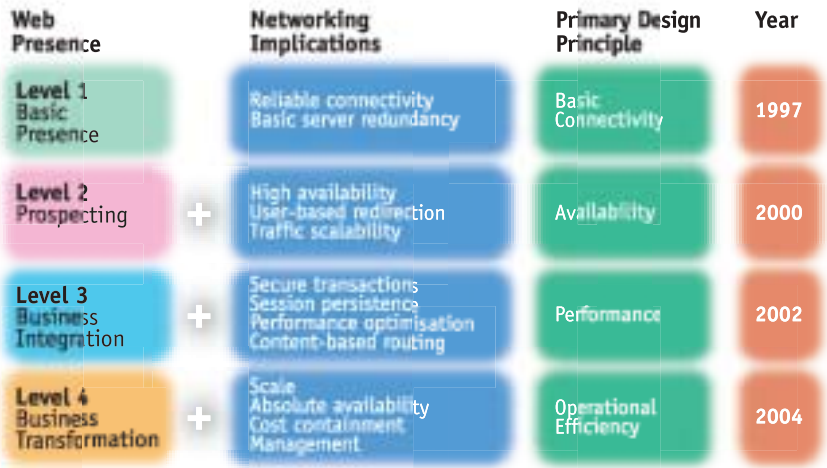

Source: Gartner Research

Many enterprises have reached Level 2 and are moving towards Levels 3 and 4 to optimise performance and achieve operational efficiencies. Application Traffic Management is key to obtaining the benefits from secure, high performance content-based routing and provides the platform to enable business transformation benefits.

According to Gartner research authored by analyst Mark Fabbi, 'Enterprises often take an overly tactical and short-term view of their content networking requirements. They apply a temporary solution to a single problem, or a relatively simple application is used as a pilot that will evolve to a more complex environment. Enterprises that consider longer-term and more comprehensive approaches to content networking, including the hierarchy of needs model, will have stronger, better-performing, longer-lasting solutions. Although upfront capital costs will likely be higher, we believe that enterprises that take a more strategic view of their Internet Data Centre requirements will decrease their total costs by at least fifteen per cent and will significantly increase the performance and reliability of their infrastructure by enabling it to more quickly meet the changing needs of their business.'

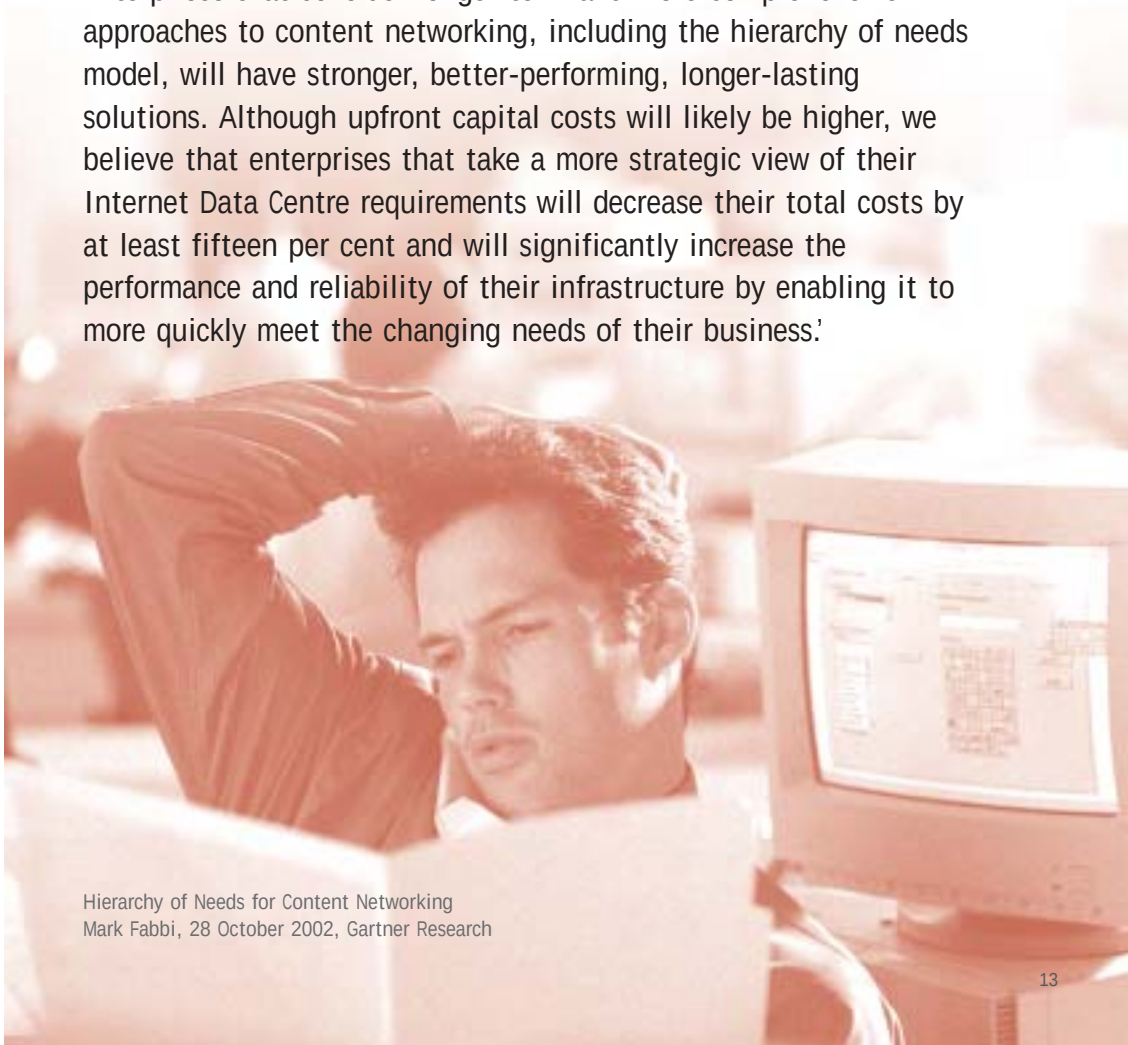

Hierarchy of Needs for Content Networking
Mark Fabbi, 28 October 2002, Gartner Research

Application Traffic Management in Action

The following scenarios demonstrate some common application problems that, prior to Application Traffic Management, could not be solved cost-effectively or implemented without significant complexity and ongoing maintenance burdens.

A Financial Institution – supporting priority customers

A US Financial Trading Network needs to scale and persist client web service traffic based on the type of traffic and its priority scheme whereby unique 'gold' customers are sent to 'gold' priority servers. Before Application Traffic Management the financial institution would have needed to architect switching and persistence functions directly into the application. A feasibility study showed this approach had a high development and maintenance cost – around $1m annually - plus an unknown time to market and approximately 30% impact on server performance, hence the need to increase computer hardware by a third.

Instead, by using Application Traffic Management technology, BIG-IP can identify clients and their individual priority status. The Universal Inspection Engine, directed by specific iRules, can direct customers to the appropriate array of servers and persist client transactions until completed. This solution has a significantly lower cost without the time-to-market uncertainties.

A Logistics Company – scaling an online parcel tracking system

A global transportation company provides an online tracking service to its customers and staff, powered by high availability databases. The majority (85%) of requests to the databases are 'read only' operations, growing at 90% per annum – typically customers enquiring on package delivery status. Prior to Application Traffic Management, in order to scale the database as business volumes grow, the only solution would be to buy expensive mid-range and higher-level servers and replication software. The replication process is very expensive and complex to maintain. In a redundant configuration both the read and the write servers would need to be equal. Thus the larger the application, the more expensive the implementation – calculated in this instance at around $750,000.

Today, however, by using a pair of BIG-IP and less expensive mid-range servers, the logistics company is able to switch and segment traffic based on whether it's a read or write operation. Read operations are sent to clusters of low-cost web servers which can be scaled quickly, while write operations are sent to mid-range servers that periodically update the read servers. In this example, cost savings in excess of $400,000 have been estimated by architecting a more efficient IT deployment model founded upon Application Traffic Management procedures.

A Large Manufacturer – securing transactions with trading partners

A large manufacturer wants application-level security for business-to-business transactions with its suppliers, but at a practical cost. Firewalls simply don't offer the level of security required between more than fifty suppliers.

A traditional solution would be to 'architect-in' the security required into its applications with all the associated development, maintenance and new hardware costs as well as limited scalability and no guarantee of the flexibility needed to support evolving applications. Such a bespoke development scenario was calculated to cost between $2.5 –3 million.

However, with Application Traffic Management technology, the manufacturer can offload security functions from the servers and applications themselves and onto BIG-IP. This offers a centralised security control, which can set-up secure channels, authenticate, authorise, enforce intrusion prevention and alleviate application attacks with full reporting capabilities. Because the BIG-IP can inspect any type of application traffic and make routing / encryption decisions based in iRules, any current or future applications can have a policy-based security procedure without any re-coding of the application. In this example, savings are calculated at around $2 million by avoiding the cost of application re-development, support and additional hardware purchase.

Application Traffic Management – The Way Ahead

Organisations are seeking innovative enhancements to their IT investments in order to improve their operational efficiency – hence building an application-aware network has now become a necessity. To help address this challenge, traffic management has come a long way from its roots in local load balancing. With the advent of Application Traffic Management, for the first time, application intelligence can be automatically applied to application data flowing across and outside the corporate network. The server, the network and application software can be built into a coherent and secure system without the prohibitive costs and inflexibility of previous bespoke developments. With Application Traffic Management, the benefits envisioned through the deployment of web services, mobile computing and web-enabled applications can be realised today. The challenge now is to apply these techniques to your business.

'We still consider F5 to be the thought leader in the market. They continue to add to a broad product offering and include a dizzying array of features, some of which competitive companies have built entire product offerings around.'

Mark Fabbi *Gartner*

About F5 Networks

F5 is the acknowledged leader in Traffic Management, and has pioneered Application Traffic Management and the patent-pending technologies that underpin this important advance. As a specialist in this area, it is working closely with many major server and enterprise application vendors to ensure that the benefits of BIG-IP are being built into industry-standard platforms. If you would like more information about these and related developments, please contact your local F5 office, where our staff would be delighted to assist you with your traffic management activities.

'We believe F5 will continue to add capabilities that are compelling to the enterprise and broaden their offering to concentrate on the natural consolidation points within the IDC.'

Mark Fabbi *Gartner*