In The Boxing Ring

IN THIS ISSUE

2.
BANDWIDTH
A discussion on bandwidth usage, concentrating on web usage and the results of recent Network Box Security Response research on such usage. We also discuss the myth of the ‘long tail’, and how the 80/20 rule can affectively be applied to keeping your Internet pipes flowing clean and fast.

3.
SYSTEM RESILIENCE
A presentation on the topic of system resilience to faults, the approach of adding servers to increase reliability, and suggestions for hardening approaches to two common Internet protocols. The topics of High Availability and Load Balancing are also introduced.

4.
DECEMBER 2009 FEATURES
As usual, we will be deploying our on-going enhancements and improvements as well as maintenance features to all NBRS-3.0 customers.

4.
DECEMBER 2009 HINT
All or nothing - how mixing responsibilities can lead to a mess and result in an unclear idea as to ‘ownership’ of an issue.

Network Box Technical News from Mark Webb-Johnson, CTO Network Box

Welcome

歡迎您來閱讀2009年12月版“In The Boxing Ring”期刊。在這個版本中，我們集中討論兩個議題——帶寬統計和通過加固網際網路協定來防止系統出現故障。

在第2頁，我們將討論Network Box安全回應團隊最近對我們的客戶進行帶寬使用的一些調研結果。儘管目前有許多安全廠商大肆吹捧關於“長尾”理論的神話，但是我們的調研結果卻證明了悠久的80/20規則往往更符合現實。

轉到第3頁，您將看到針對防範系統故障的專題介紹，特別是針對常用的互聯網協定中的兩種核心協定進行加固的一些專業建議。

打開第4頁，按照慣例將提供關於月度總結和使用技巧的小提示。

和以往一樣，如果您有任何的回饋，意見或者建議，我們都歡迎您隨時提出來。您也可以通過發送郵件到我們的郵件列表： nbhq@network-box.com 聯繫我們。或者當您下次在香港市區的話來隨時來我公司辦公室進行參觀指導。

您也可以通過加入或訂閱我們的安全響應Twitter 和我們保持聯繫，網址是：

twitter.com/networkboxhq

Mark Webb-Johnson
CTO, Network Box Corporation
December 2009

NETWORK BOX

Bandwidth Usage
Network Box Survey Results

A typical Network Box customer sees his bandwidth divided into approximately 95% used for email and web browsing, with only 5% (or sometimes less) being used by other traffic. Customers who don't impose outbound policy restrictions (LAN->NET without restriction) will also see a significant amount of file sharing activity. In the rest of this article, we concentrate on the web traffic portion of the problem.

Recently, Network Box Security Response conducted a statistical survey of bandwidth usage of our customer base, concentrating on web activity for the past four months, to clearly benchmark our performance in web content filtering - the largest such survey of its kind.

While companies such as Alexa benchmark the Internet as a whole, this research was targeted specifically at our customer usage, and went beyond the Alexa 'hits' approach (ie; number of visits per web site) to include 'size' data (ie; size of data transferred from the web site).

The results were astounding. While 29% of web requests were to the top 20 sites (ranked by number of hits), 37.9% of bandwidth was used by just the top 20 sites (ranked by bandwidth used).

Facebook (facebook.com and fbcdn.net) was responsible for 5.8% of the URLs going through our boxes, and 4.4% of the bandwidth. But, that was dwarfed by YouTube (at 7.8% of bandwidth). Network Boxes processed (and our customers paid bandwidth for) over 21TB of YouTube videos in just four months.

Even Microsoft Windows Update (windowsupdate.com) at 0.9% of URLs and 3.8% of bandwidth stands out. It is amazing to think that much bandwidth is required just to fix software problems in Windows. Both Symantec and NAI (both ant-virus security companies) also appear in the top 20 list of bandwidth users.

In total, our customers visited 4.4 million web sites using over 19 billion URLs, consuming 280TB of bandwidth.

Top 20 Web Sites (by hits)		
	Site	%Hits
#1	google.com	4.1%
#2	fbcdn.net	3.4%
#3	yimg.com	3.1%
#4	yahoo.com	2.7%
#5	facebook.com	2.4%
#6	doubleclick.net	1.7%
#7	msn.com	1.5%
#8	google-analytics.com	1.1%
#9	news.com.au	1.0%
#10	nextmedia.com	1.0%
#11	microsoft.com	0.9%
#12	bbc.co.uk	0.9%
#13	windowsupdate.com	0.9%
#14	theage.com.au	0.7%
#15	live.com	0.7%
#16	gstatic.com	0.6%
#17	iitech.dk	0.6%
#18	googlesyndication.com	0.6%
#19	on.cc	0.6%
#20	sinais.cn	0.5%

Top 20 Web Sites (by size)		
	Site	%Size
#1	youtube.com	7.8%
#2	windowsupdate.com	3.8%
#3	yimg.com	2.8%
#4	google.com	2.7%
#5	fbcdn.net	2.3%
#6	facebook.com	2.1%
#7	microsoft.com	2.0%
#8	rapidshare.com	1.9%
#9	apple.com	1.8%
#10	symantecliveupdate.com	1.6%
#11	yahoo.com	1.6%
#12	adobe.com	1.5%
#13	megaupload.com	1.2%
#14	googlevideo.com	0.8%
#15	nextmedia.com	0.8%
#16	mediafire.com	0.8%
#17	llnwd.net	0.6%
#18	msn.com	0.6%
#19	2mdn.net	0.6%
#20	nai.com	0.6%

Looking at just the top 1,000 sites, we find 74.6% of URL requests, and 76.3% of bandwidth are used. For the top 100,000 sites, the figures increase to 97.1% and 96.9% respectively. **97% of the web traffic (and bandwidth) goes to the top 2% (100,000 out of 4.4 million) of sites.** Certainly, the 'long tail' on the chart is impressive, but if you want to control bandwidth usage, and make an impact on the load of web proxy traffic on your network, then targeting these top 2% of sites is going to make by far the biggest impact.

Regarding categorization of web sites; of the top 100,000 websites, SurfControl categorizes 78.6%, and Network Box (including all engines) categorizes 88.2% (including 1,692 as undesirable). The difference is the result of the Network Box uncategorised URL feedback loop (as presented in the July 2008 "In The Boxing Ring" newsletter), as well as our partnership with other categorisation list providers (in addition to SurfControl).

Network Box finds that the requirements for the 'fat head' and 'long tail' of the URL chart are dramatically different. At the head, the accurate categorization of all sites is of upmost importance. In the tail, the categorization of malicious and undesirable websites becomes most important (while categorisation of non-malicious sites is relatively unimportant).

The Network Box approach is to apply full database-based categorisation at the head, and rely on more dynamic categorisation of the much smaller number of undesirable sites in the tail. Whilst we are currently meeting our goals, we continue to work on improvements to our core technology.

構建更具彈性化的系統
通過加固互聯網協定積極防範系統失敗

讓我們從破解一個眾所周知的迷思開始：增加伺服器的數量並不會減少故障發生的次數。相反，它增加了故障發生的次數。例如，假設一個特定類型的伺服器每12個月會出現一次故障。然後，將在兩個伺服器一起計算，平均每6個月就會出現一次故障。故障發生的可能性會隨著部署伺服器數量的增加一起增加。

然而，如果你的伺服器可以彼此從發生故障的其他伺服器接管過去，您就以通過部署更多的伺服器來減少在同一時間內所有伺服器都發生故障的可能性，從而提高整體服務的可用性。

問題是，當然——細節問題，部署當出現故障時彼此可以順利接管的服務，這些故障分爲微小到可以忽略的、互有依存關係的、大規模的和服務本身的。本文的其餘部分將詳細討論兩種常用的互聯網協議以及如何對它們進行加固以防止故障發生。

功能變數名稱系統（DNS）

DNS系統運行在TCP或UDP埠53上。用戶端配置成使用多個DNS伺服器之一——如果一台伺服器沒有回應，然後在用戶端切換到下一個。

因此，添加多個DNS伺服器列表會讓使用變得更加可靠。但是，如果列表中的DNS伺服器失敗，切換時間往往是7.5秒（或以上），接著之後經過它的所有查詢可能會受到不利影響。如會造成一些任務如網頁流覽速度明顯放慢。

另一個問題是，目前DNS系統，限制了只能有效地使用不超過512位元組的UDP資料包。增加DNS伺服器的數量可能會造成回應資料包突破512位元組的限制，進而將連接切換到TCP/53（這種方式更慢，而且經常被防火牆阻止）。

Network Box建議您使用一個合理數量（1到3）且高度可靠的DNS伺服器，而不是幾十個不可靠的DNS伺服器。

您還應該注意配置您的TTL（存活時間）和SOA（權威設定）序列號的正確，以便所有的DNS伺服器配置正確和及時同步。

SMTP郵件

SMTP郵件系統運行在TCP埠25上。用戶端配置成使用一個出站SMTP郵件伺服器。

出站SMTP的情況 相對較簡單。Network Box建議，如果你要部署多個SMTP郵件伺服器（用於出站郵件），您需要指定名稱伺服器（使用較短TTL的DNS），而不是IP位址——這將允許您更快 地切換到備用伺服器。如果您有一個高可用性配置，您也可以通過虛擬IP的使用，在伺服器發生故障時實現無縫切換。

入站SMTP的情況較爲複雜。您 需要在DNS系統中發佈MX記錄，登記那些可以接受您的功能變數名稱電子郵件的SMTP伺服器列表。注意在DNS配置中不要超過512位元組的UDP資料包這個限制。MX記錄包含優先清單，郵件伺服器應該連接到您最優先的郵件伺服器——但垃圾郵件發送者經常違反本規則，他們喜歡連接到較低優先順序的郵件伺服器。

對於入站SMTP，Network Box強烈建議你列出一條或兩條MX記錄，並使用不同的優先順序別。您的Network Box也可以被配置成備份MX，以便讓郵件佇列存在其中推遲傳遞給您的伺服器（當您的郵件伺服器出現故障時）。

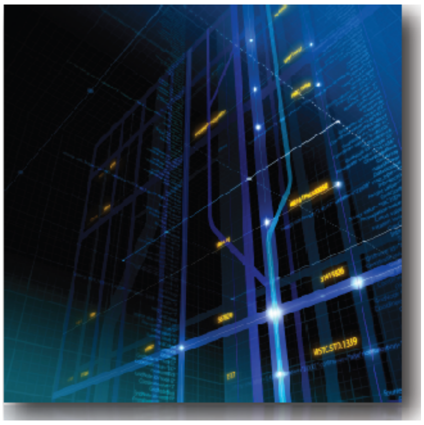

硬體彈性化

對於在辦公環境中的使用，Network Box建議您採納“保持簡約”的格言。使用複雜的負載分擔似乎是個好主意，但大多數辦公環境能夠容忍一些停機時間——保持簡約可以避免給您的網路引入大量新的問題。

對於那些需要全天候不間斷可用性的客戶，Network Box可以爲您提供既高可用性和負載平衡。這些是通過配置一個虛擬IP來實現的，這個虛擬IP可以由兩個（或多個）物理機器共同分享——只有其中的一台機 器（主）持有虛擬IP位址。

如果主機器失敗，一個虛擬的選舉將馬上舉行，最高優先順序的備份機將接管虛擬IP位址和主機角色（當原來的主機器恢復回來時交回 虛擬IP和相關角色）。

硬體彈性化、高可用性和負載平衡的主題是非常複雜的。

設計這樣的系統需要對核心的互聯網協議是如何工作的有一個非常深刻的理解，同時要用預見的能力，並允分計畫應對不同類型的失敗。

如果您需要這方面的更多幫助，請聯繫當地的Network Box網路安全運維中心。

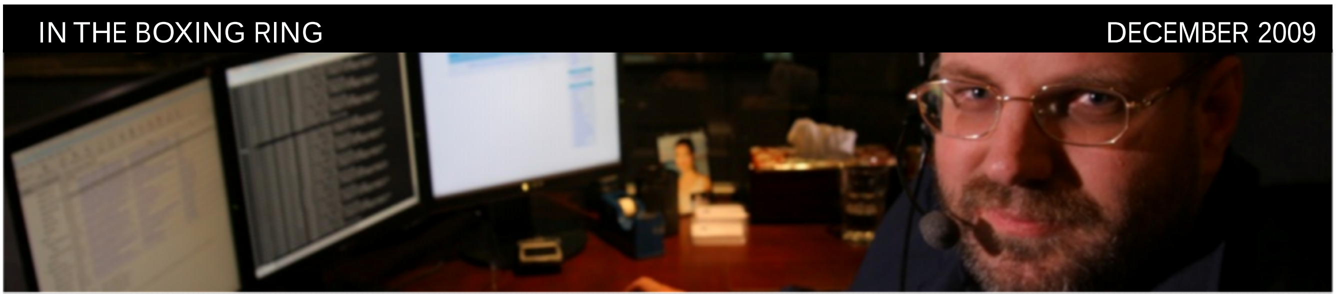

December 2009 Features

On Tuesday, 1st December 2009, Network Box will release our patch Tuesday set of enhancements and fixes. The regional NOCs will be conducting the rollouts of the new functionality in a phased manner over the next 7 days. This month, these include:

- Enhancements to the POP3 Acceleration system to improve performance for large numbers of mailboxes, and introduce a new mechanism to timeout connections to non-responsive POP3 servers.
- A fix to the policy categorisation caching system, to better handle the case where a domain and it's subpath are categorised differently, in particular when using complex categorizations using regular expressions.
- Enhancements have been made to the Network Box Signature PUSH system to improve performance and reliability when delivering simultaneous signature updates.
- Add support for challenge-response to the envelope verification system used for mail scanning.
- Enhancements to the LDAP integration system, to support more complex directory environments.

In most cases, the above changes should not impact running services or require a device restart. However, in some cases (depending on configuration), a device restart may be required. Your local NOC will contact you to arrange this if necessary.

Should you need any further information on any of the above, please contact your local NOC. They will be arranging deployment and liaison.

December 2009 Hint

我們收到的常見要求之一是各種形式的“混合”控管。問題是，用這樣的系統，職責的混合可能導致部署的混亂，進面導致問題的“所有權”沒能明確界定。

舉例來講，我們經常看到（與第3頁的DNS彈性有關）Network Box將使用哪種DNS伺服器的專門設置。每台Network Box都包括一個完全符合標準，高可靠，可配置的並且使用BIND互聯標準代碼的DNS伺服器。我們經常建議客戶將他們的網路配置成使用Network Box作爲他們的DNS伺服器。

然而，經常有客戶請求我們使用客戶自行配置的DNS伺服器，再加上Network Box上的DNS伺服器。正確的做法是採用如下方法之一：

(a) 轉發客戶功能變數名稱（基於由客戶提供功能變數名稱列表）解析請求給客戶的DNS伺服器，但允許Network Box的DNS伺服器去解析除它之外的其他功能變數名稱，或

(b) 配置Network Box完全使用客戶的DNS伺服器。

不正確的方法是僅將客戶的DNS伺服器作爲其一加入到Network Box用的DNS伺服器列表中——因爲這樣做當有DNS問題發生時可能造成混亂（必須花費時間和精力確定所使用的眾多DNS伺服器中到底是哪台出了問題）。

Mark Webb-Johnson,
CTO, Network Box Corporation

NOVEMBER 2009 NUMBERS

Key Metric)	#	% difference (since last month)
PUSH Updates	1,246	-8.4
Signatures Released	236,109	+7.1
Firewall Blocks (/box)	642,767	+4.8
IDP Blocks (/box)	210,821	+10.7
Spams (/box)	64,902	-0.5
Malware (/box)	2,905	-28.4
URL Blocks (/box)	100,429	-5.4
URL Visits (/box)	3,175,339	+7.7

NEWSLETTER STAFF

Mark Webb-Johnson
Editor

Michael Gazeley
Jason Law
Nick Jones
Production Support

Network Box Australia
Network Box Hong Kong
Network Box UK
Contributors

SUBSCRIPTION

Network Box Corporation
nbhq@network-box.com
or via mail at:

Network Box Corporation
l6th Floor, Metro Loft,
38 Kwai Hei Street,
Kwai Chung, Hong Kong

Tel: +852 2736-2078
Fax: +852 2736-2778
www.network-box.com