me
Nathaniel Cohen
ncohen@ucsd.edu
github: @breathe

Programmer/Analyst with 7 years (yikes!) experience as a sea-going science-supporting IT technician

philosophie:

- glossary

- is there a bandwidth capacity deficit??

- managing bandwidth is about finding ways to *equitably* reduce demand for a limited resource while meeting primary objectives

technik:

- captive portals

- data quotas

- deployment of cyberoam and peplink appliances on New Horizon/Revelle

quick terminology overview

flow, packet flow — A sequence of packets observed to traverse the network in one direction from a source to a destination by a particular intermediate device within some span of time.

an artificial *logical* equivalent (in a hand-wavy way) in packet switched networks to the concept of a connection in non-packetized networks

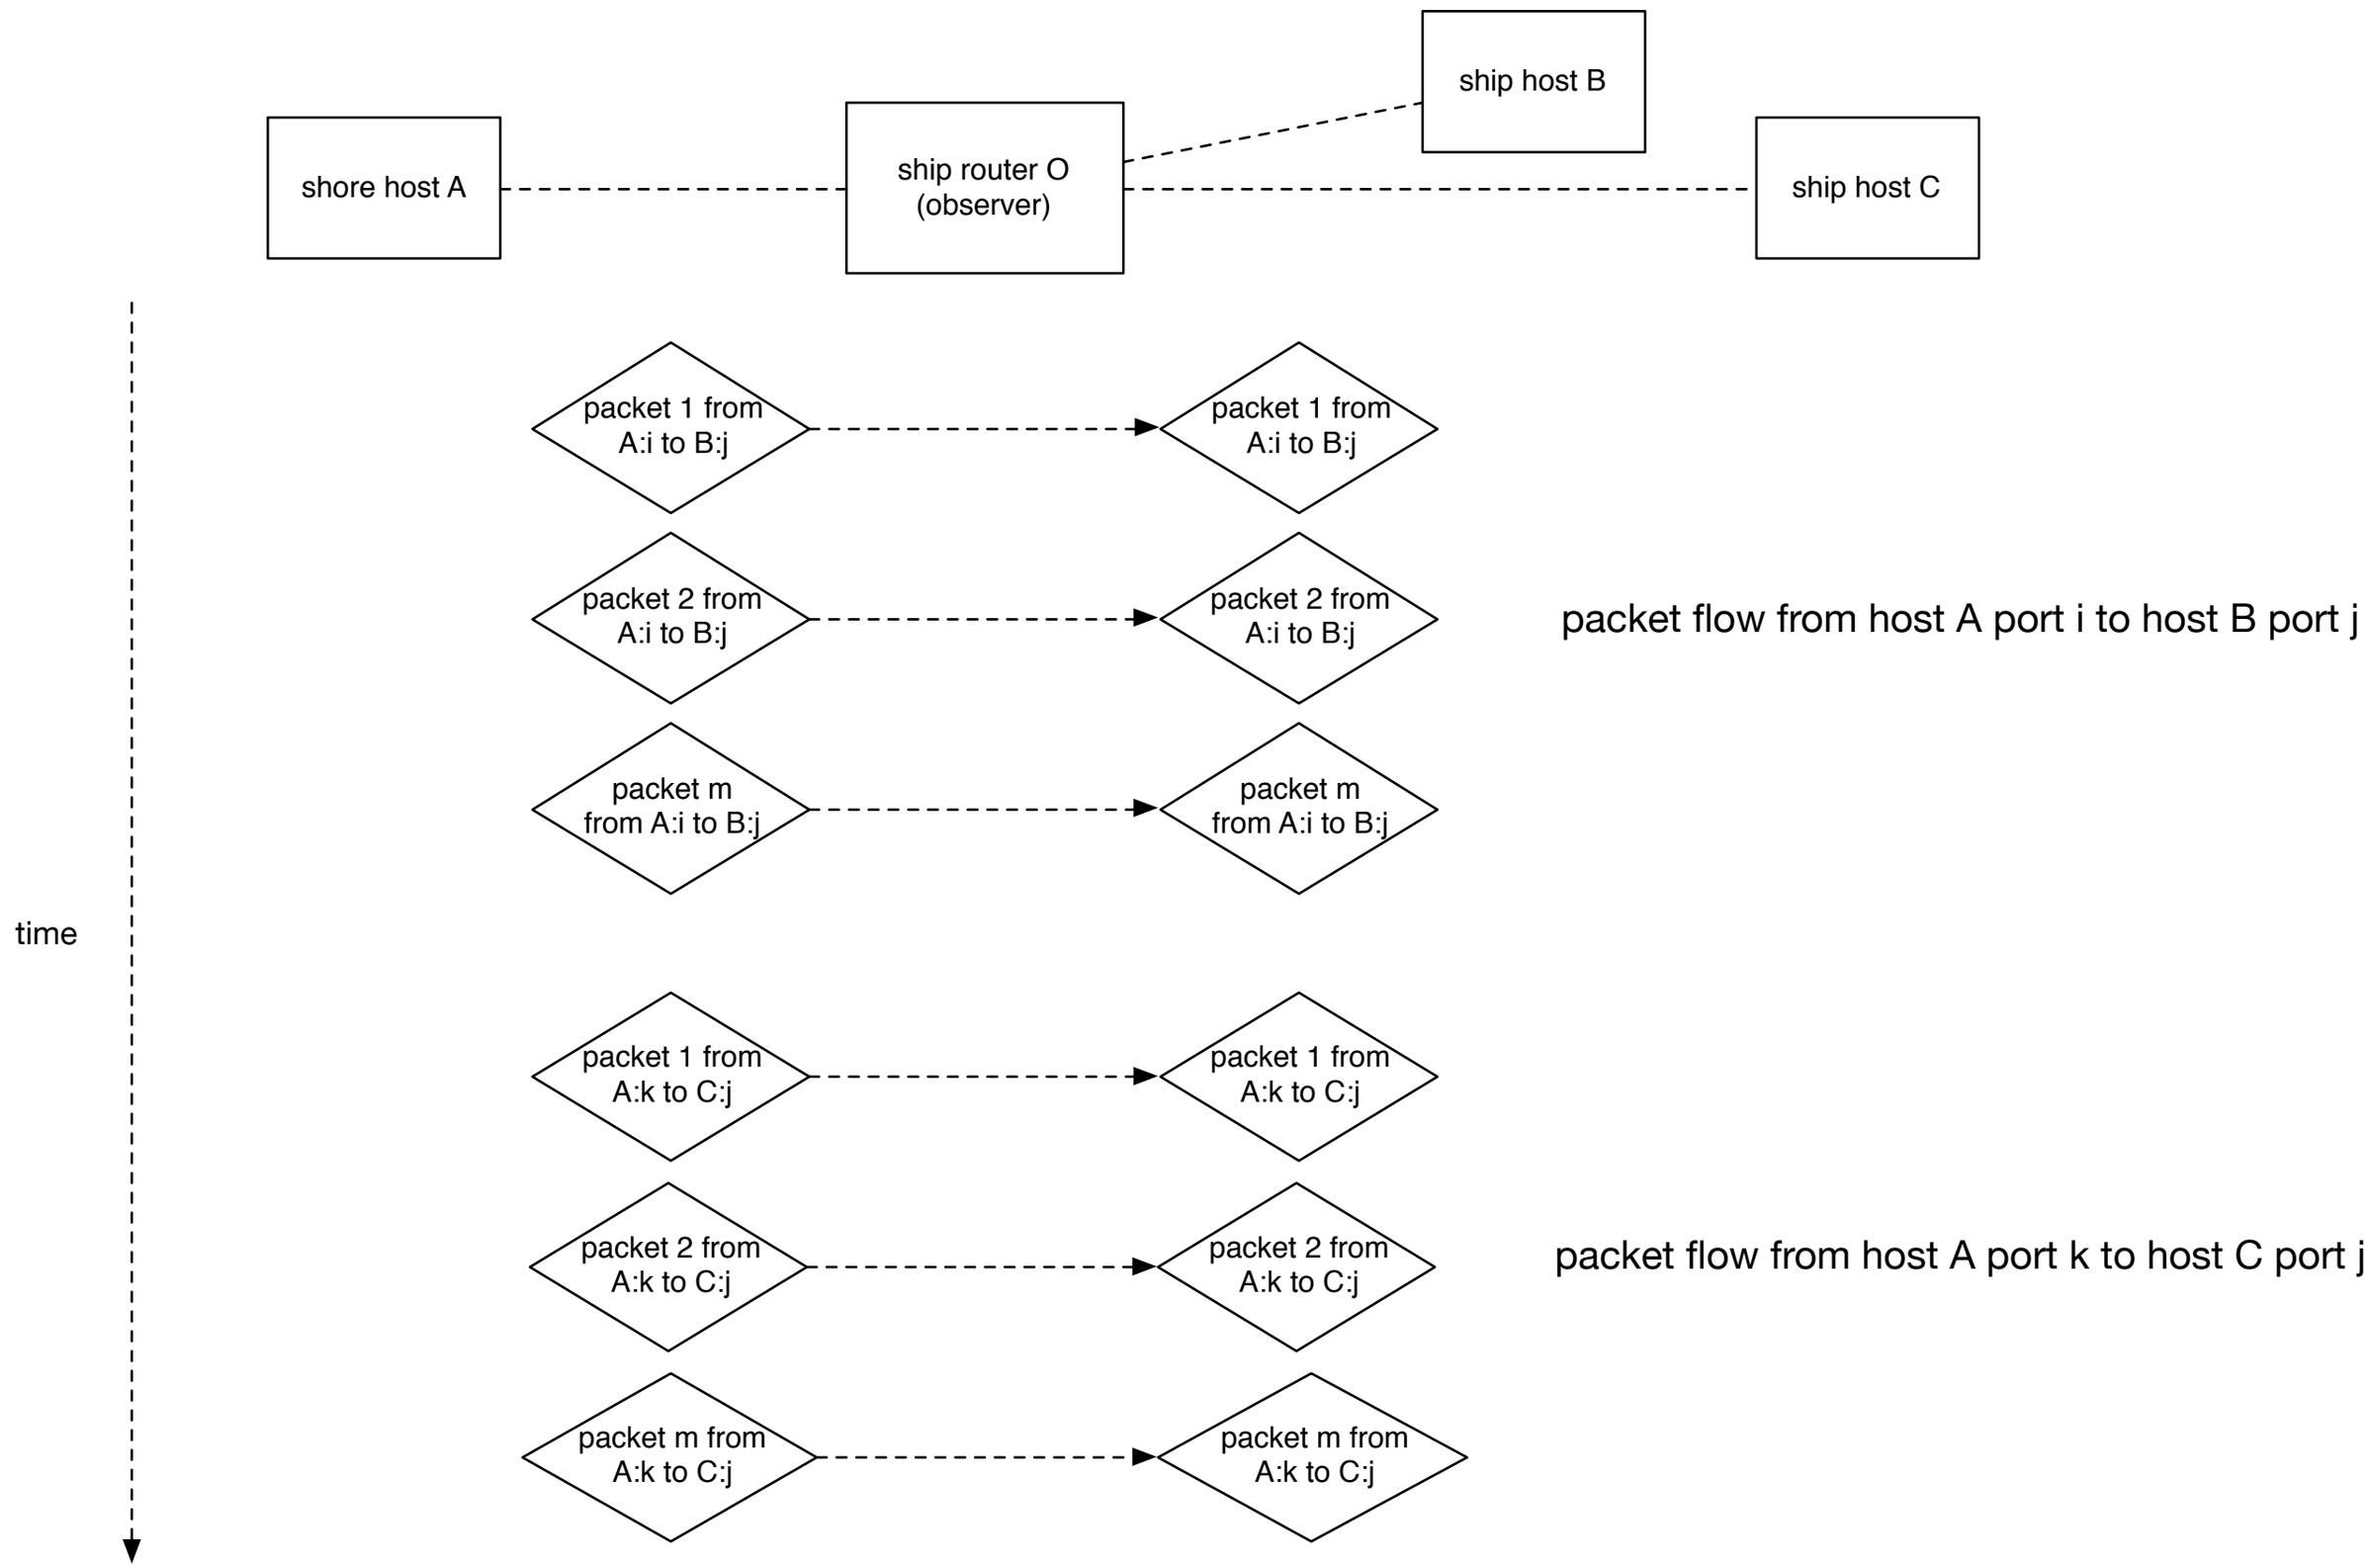

session — Not strictly defined, but its the concept of a flow or bi-directional pair of flows along with a lifetime determined by an observer O, which would most typically be a nat router or a firewall.

example: when a nat router/firewall allows an outgoing connection, an entry is added to the nat-router's 'nat table' or firewall session table telling the device to 'allow packets associated with this connection' — the entry includes some kind of lifetime measurement saying for how long this 'session' should stay alive or more generally saying how to determine when the session is no longer alive

"User internet expectations have grown enormously since internet was introduced in early 2000s"

THEN

Internet still referred to as 'information superhighway'. Before the internet was realized to be a series of tubes.

- The majority of sailing personnel did not have or bring laptops with them to vessels
- very few users amongst crew and science party were conditioned to use the internet regularly
- some users almost never used the internet even when ashore — and virtually no one used it continuously throughout their normal daily existence
- software still commonly updated by cd/dvd

Large uses of data services by byte and by user interest

- video delivery needs (when required)

- emails/messaging — typically low-bandwidth, almost no one had digital cameras. Email tended toward small plain text messages

- scientific data transfer needs often represented a large part of the total demand for the shipboard network's internet capacity

- web browsing content of the era was still being optimized (to a certain degree) for bandwidth efficiency as a general rule

- browsing the web generally meant reading text with maybe some low-res thumbnail images which could be clicked to download higher-res versions. Websites were static entities — they were more cacheable then than now. Once fully downloaded connections to the site ended and javascript rarely lead to the download of additional content without user interaction.

- imagery deserving the label 'high-res' often meant a million pixels

NOW

Post-mobile revolution, web 2.0, retina displays, smart phones, 'the cloud', and ubiquitous high speed LTE networks …

- *ubiquitous internet enabled devices* — the majority of science party/crew will come aboard with 2 or more internet capable devices
- the majority of users amongst science party and crew are now conditioned by very deeply ingrained habit to use the internet regularly
- everyone uses the internet for personal business, for communication, for facebook — regularly and often
- overall consumer software ecosystem has evolved to 'expect' that access to a high speed network is available or that it will at least become available very soon …

Large uses of data services by byte and by user interest

- video delivery needs (when required)

- email/messaging — everyone has digital camera(s). text communication mechanisms commonly used as a medium for exchanging large digital artifacts (photos, videos) as they are for text only by many users

- scientific data transfer needs represent a tiny fraction of the demand for the shipboard network capacity given the explosive growth in other demands

- web content optimized increasingly for LTE, retina assets increasingly the norm even for mobile sites. Many javascript heavy sites generate network load 'continuously' even when a user is not active, pre-caching content the user is likely to look at 'soon' even if they never do

'high-res' label for images likely to mean 3-10 million pixels or more

- Average size of a web page has more than tripled since 2012 — compare to early 2000s …

- background software activity is ubiquitous and increasingly moving beyond the scope of most user's capacity to control without learning obscure things

Headwind summary:

- internet capacity has doubled approximately once since early 2000s from 256kbps shared amongst up to 4/5 ships, to 512kbps shared. Compared to thousands of times greater growth in shore network speeds.

- consumer devices will continue to chase after future increases in shore internet capacity growth with features reliant on internet access

- *more* people with *more* internet needs from *more* devices

"The bandwidth deficit of shipboard environments is real"

demo

Most popular website in 2014

www.facebook.com

Most popular website in 2003

https://web.archive.org/web/20031117213651/http://yahoo.com

Optimizing user experience when faced with a bandwidth deficit

newhorizon 5-minute bandwidth plot

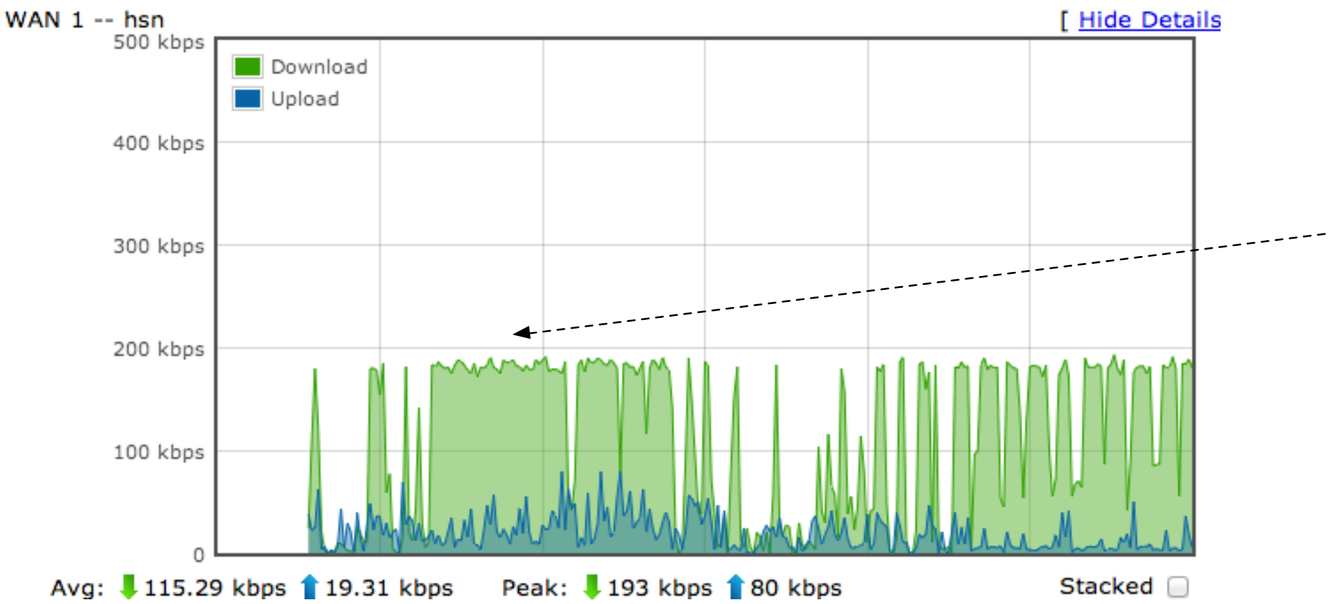

sustained peak during a period of time when only a single tcp flow was active on the network — notice the peak throughput

The purpose of this plot is simply to illustrate a simple truth about the performance of existing ship internet services:

Any single active tcp talker from the ship talking to virtually any service on shore will hit a bottleneck at the satellite hop. One device online with one active tcp session will saturate the link — anyone on shore can essentially always talk to us as fast as we can hear it.

The throughput of the ship internet is essentially a step function

the network is either totally pegged since with even one user attempting one network transaction - our meager satellite link's capacity will be reached, or there's nothing going on at all.

Under the normal case where there are multiple simultaneous active tcp flows, the throughput of each flow (assuming gear is configured properly) will generally converge toward an 'even' division of the available satellite capacity. N competing active tcp talkers implies each talker gets roughly 1/N of the link capacity**.

**Except in pathological cases when congestion collapse occurs — when N is too large relative to the link capacity, tcp will fail to converge on a stable split of the link capacity — the goodput achieved by the network in this case can degrade very *severely*

Some (mostly universal) truisms:

(1) the best possible user experience is when there are as few active flows competing at the bottleneck as possible.

(2) since there are many competing purposes for the link we must allow multiple people to use the link at the same time — competition for this limited resource is inevitable.

Many techniques exist to arbitrate network resource contention. Some examples include qos (network attempts to give some traffic priority over others) and systems like quotas (network enforces limits on the amount of traffic from individual talkers in some amount of time).

As a general rule under (with or without a qos scheme): the fewer the number the number of active talkers on the network at any given time, the better the achieved experience per user will be at that time.

Minimizing the average number of talkers ensures the highest average throughput per flow and delivers the best user experience achievable at that link capacity.

Bandwidth management (as a distinct concept from QoS) is about trying to reduce demand in order to raise the average throughput per flow.

Reducing the average number of simultaneously active users/flows will deliver a better observed network experience per user since there is a greater probability that there will be fewer other flows to compete with during that user's activities.

"managing bandwidth usage to optimize for user experience means reducing the average number of simultaneously active users/flows while allowing users to accomplish their highest priority objectives"

principles for conserving bandwidth

(1) each user needs to be motivated to contribute to and receive benefit from the bandwidth conservation process (to avoid tragedy of the commons)

(2) user's need some kind of feedback that they are capable of and motivated to respond to. If a user doesn't want or is unable to change their practices then at minimum the impact they have on the rest of the user's network experience should be bounded.

concepts for optimizing user experience

(1) every effort should be made to prioritize interactive use (user in front of computer) over (often unintentional) background use
(2) the fewer simultaneous flows on the network, the better the experience delivered to the user per flow.

Optimizing for user experience means achieving the desired user objectives with the fewest number of simultaneous flows — this implies the user achieves their objectives having encounteried the smallest amount of contention with other users while doing so …

Arguments for a Captive Portal + Daily Bandwidth Quota System:

- motivate users to adapt their and their device's behavior

- provide feedback to users to help/force them to learn about the effects of their devices' activities

- reward users who find efficient ways to accomplish their goals

- disconnect users/devices from the link for periods of the day after a quota is reached.
This implies a reduction in the average number of active device's vs the non-quota'd environment

- provide a straightforward means to take bounded advantage of non-all-you-can-eat ISP's with fixed/predictable costs (fbb … global express?)

- find and limit the effect of 'bandwidth hogs' automatically

- imply a mechanism whereby the user explicitly signals to the network 'I wish to be using the internet right now from this device' (login/captive portal).
This improves discrimination between 'user in front of computer' network activity vs 'random background' network activity

- manage users rather than devices. Restrict users to one device at a time. Users often prefer to use different devices for different tasks — allowing this sometimes both improves the user experience and reduces bandwidth usage (example: mobile app vs desktop website)

SIO recently re-organized shipboard networks on New Horizon and Revelle.

Goals:

Deploy a shipboard internet solution that

1. is easy to manage

2. is robust against the spectrum of different layer-2/layer-3 network problems

3. provides seamless as reasonably possible internet connectivity via multiple different internet links (network failover)

4. allows for the possibility to improve network performance in a simple manner by taking simultaneous advantage of multiple different wans (network load-distribution)

5. allows for problems to be easy to identify and diagnose (without requiring full-time cisco ios engineers on staff)

6. is cleanly architected so that the individual parts of the network are easy to change and evolve

7. exposes useful and approachable information regarding network operation to shipboard tech-type personnel

8. uses the simplest possible interfaces across organizational boundaries with strong decoupling between separately managed organizations (eg UCSD IT managers, STS, HSN, and any other internet providers have flexibility to make any needed changes without triggering failures due to hidden/ unrealized dependencies between the administrative domains)

The tools we picked:

Cyberoam 25iNG

'Next generation' security appliance

Provides captive portal+layer-7 firewall with url filtering capabilities.
Flexible management web-ui + reporting features

~$500
~$100/year for subscription to traffic categorization databases and updates

Peplink Balance 510

Multi-Wan Router Appliance

Provides session-based multi-wan features + speedfusion
Nice, easy, management and status monitoring features

~$2500 for 5-port model

Install: The peplink balance is connected to each of the ship's WANs (hsn, fbb, 4g modem, shore cable, etc etc …). The cyberoam connects the peplink device to the ship network.

he peplink device/configuration decides which path a packet leaving the ship should take. The peplink also maintains a speedfusion tunnel back to shore for remote connections into the ship
The cyberoam captive-portal/firewall sits between the peplink and the ship network and decides what packets are allowed on/off the ship network.

Network Overview

Decision process for traffic flows initiated off-ship

Device map

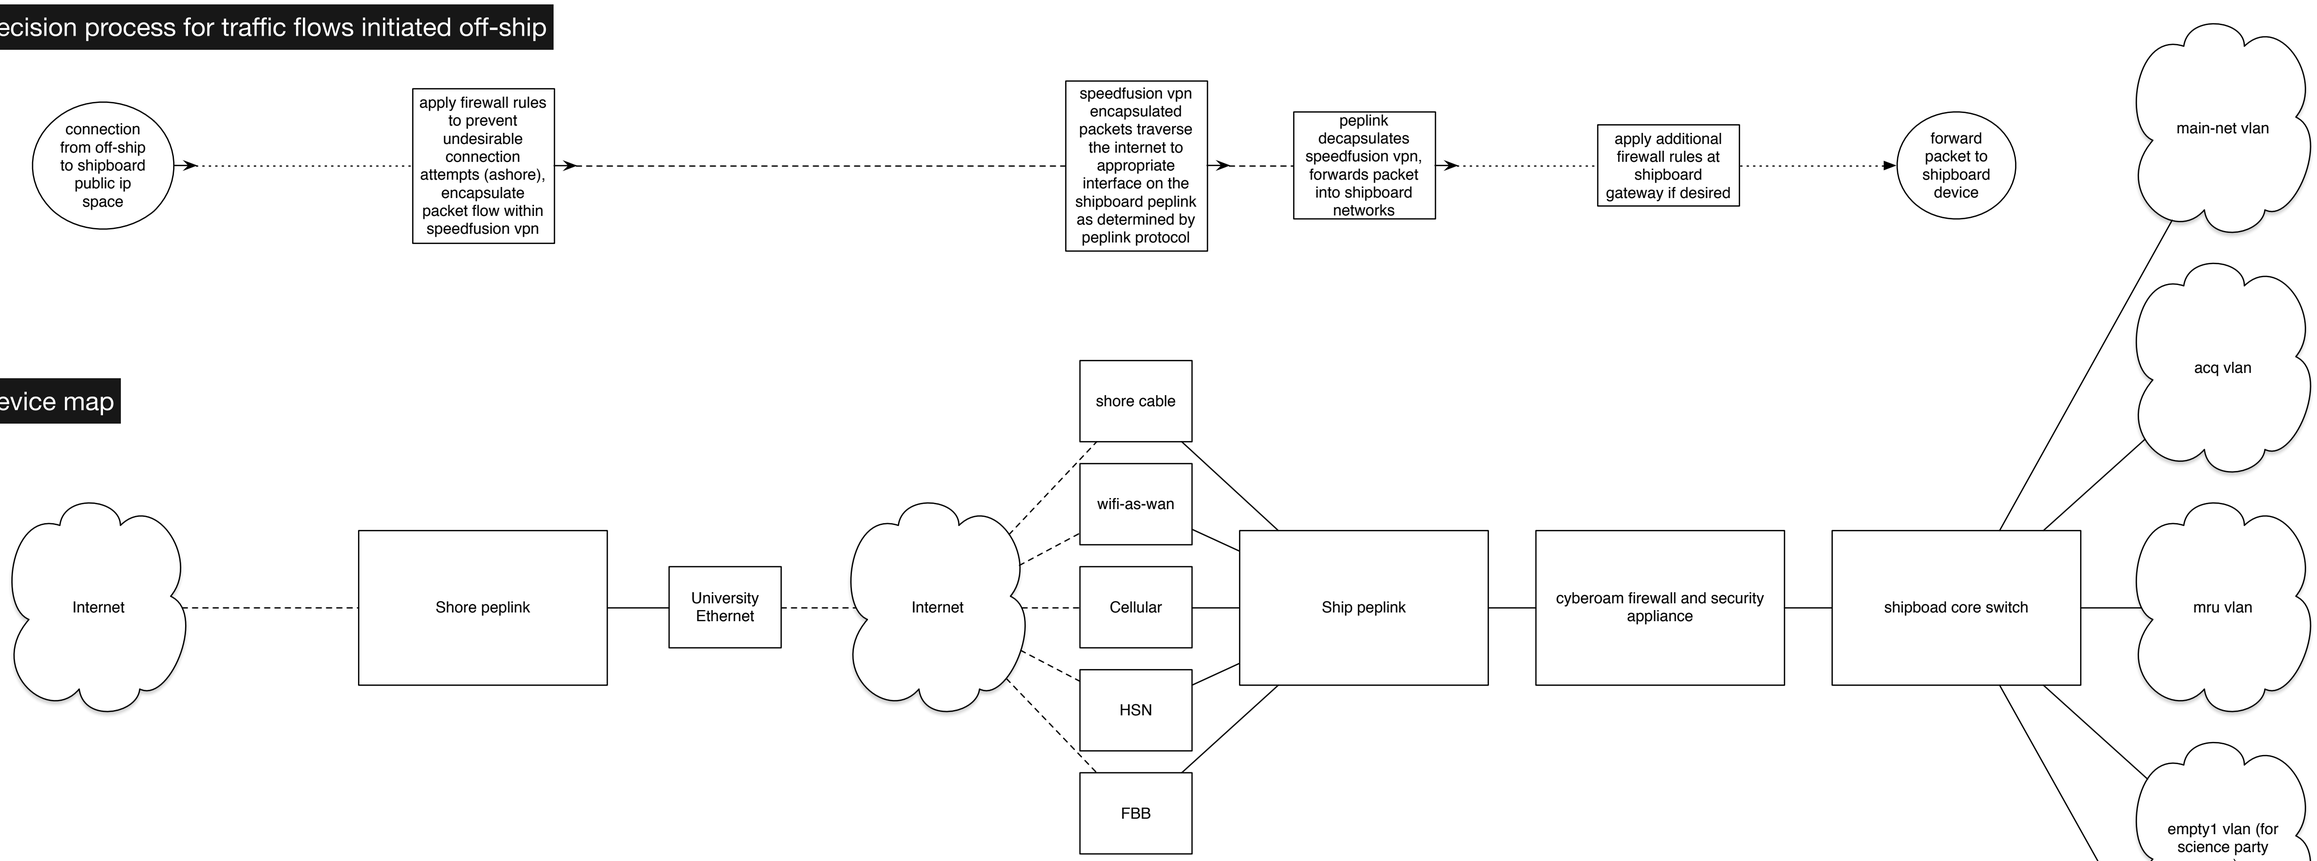

Decision process for traffic flows initiated on-ship

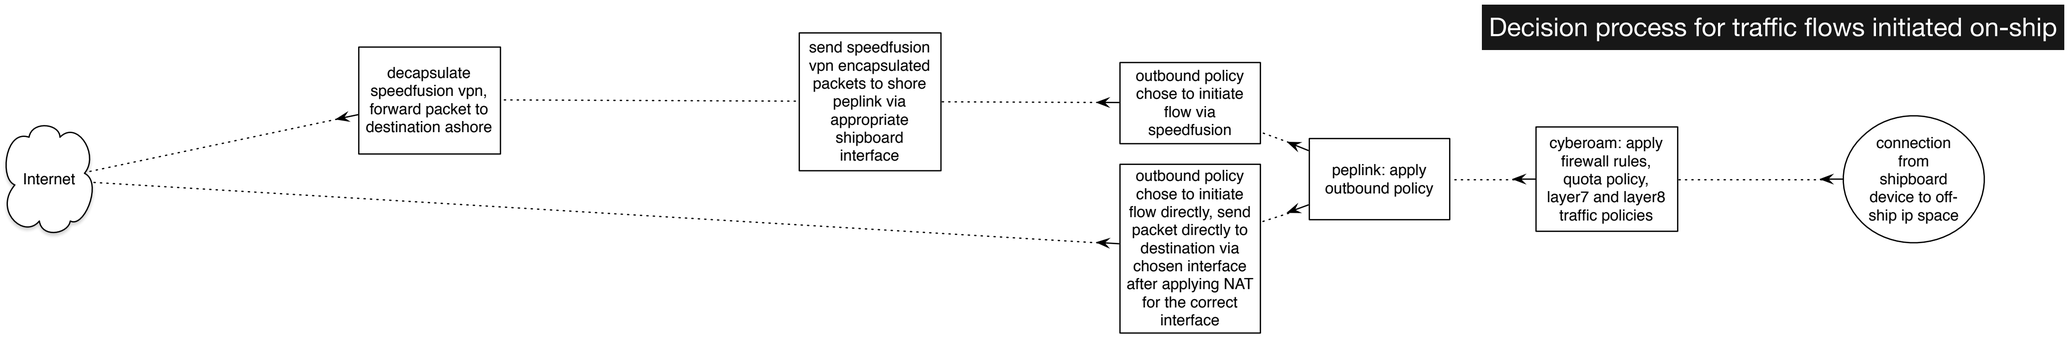

There are three diagrams shown above:

- The middle diagram shows the network equipment and layer-2/layer-3 network connections between
- the top diagram shows how the decision process works for traffic **initiated** off-ship destined for the ship's public ip space
- the bottom diagram shows how the decision process works for traffic **initiated** on-ship destined for off-ship ip space

The top and bottom diagrams are vertically aligned with the middle diagram to show which equipment implements which piece of the traffic flow logic

Note: Not depicted in the diagram above — campus network administrators maintain a simple static route to the public ip space of each ship via the shore peplink. This makes our interface with campus administrators incredibly simple 'hey can you guys publish a route to the 'revelle' ip space as being via the shore peplink device at ip x.x.x.x. Its a very simple interface to coordinate across groups -- and easy for everyone to maintain/understand.

A packet destined for the ship ip space which reaches the shore peplink is forwarded to the ship over the appropriate speedfusion tunnel. This is how we make the shipboard ip space accessible. Our shipboard public ip is accessible via all wans (and any wan we may add in the future) without doing any bgp or running any kind of dynamic routing protocols whatsoever

What is speedfusion vpn?

Speedfusion is a multi-wan capable vpn protocol implemented on peplink devices — it establishes a vpn tunnel between two peplink devices via any number of WANs on each end

VPN traffic over the speedfusion tunnel can be

1. bonded — meaning traffic can flow over multiple ip links simultaneously resulting in higher maximum throughput than achievable via any one WAN

2. used to provide a vpn link which is extremely robust against network failures. Any constituent link belonging to the speedfusion tunnel can come and go and the traffic will keep flowing as long as at least one link stays up

From the administrator perspective, it works like this — each wan which is part of the speedfusion tunnel is assigned a priority number, traffic flows via the highest priority available wan — if there is more than one wan with the same priority then the links are 'bonded' and the traffic can flow via both at the same time.

Do you use the bonding feature of speedfusion to improve network performance beyond the capacity available from one WAN?

No, in our setup each wan has a different speedfusion priority — we use speedfusion for traffic which we want to flow without interruption despite the comings and goings of the underlying wans. One big reason not to bond is that achievable throughput over a bonded link tends to suffer when the WANs involved have drastically different latencies.

We can and do still gain benefit from using multiple wans at the same time — we use a different feature of the peplinks for this — called outbound policy. The outbound policy provides load-distribution algorithms which ,do , work well despite differences in WAN latencies (more info later).

100% Packet Level WAN Bonding

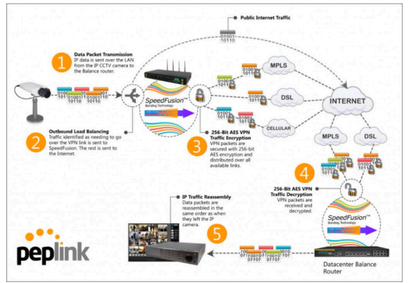

What about creating and managing the speedfusion VPN — is that hard?

simple requirements
1. shore-node needs public ip address
2. the different wans connected to the ship peplink can have any mix of static, dynamic, publicly-routable, and privately-routable addresses

easy to observe what's happening with the speedfusion state

Setup procedure:
1. define a speedfusion profile on shore node for each ship
2. on ship peplink, define a speedfusion profile with name matching the name of a speedfusion profile created on the shore node. Point the profile configuration to the ip address of the shore node
3. Done

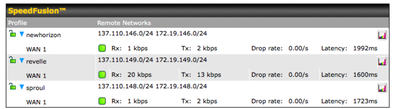

Session level vs packet level multi-wan load distribution/failover techniques

Speedfusion is a packet-level bonding technique meaning each packet in a flow is eligible to traverse the internet via any of the available WANs. Should one of the WANs go down, any WAN can continue moving packets associated with any established flow without* interruption

Another approach with a simpler implementation is "session-level" load-distribution. Session level load-distribution means each outbound network "session" is established via a particular WAN at the time the 'connection' is established. The decision about which WAN to use for a given session is made and fixed at the time the connection is established. If one of the WANs becomes unavailable, any sessions established via that WAN will be interrupted — but can be restarted easily via another WAN.

Session level load-distribution is simpler while packet-level load-distribution techniques can induce negative performance effects particularly in high-latency, low-bandwidth environments. The downside of session-level load-distribution is that connections made via a specific WAN are interrupted when that WAN becomes unavailable. The downside of packet-level load distribution techniques (of which there are a variety but generally all share these characteristics) are complexity and potential for performance impacts depending on WAN characteristics.

How hard is it to define the peplink outbound policy?

Easy - even when specifying complex behaviors

Service	Algorithm	Source	Destination	Protocol / Port	
C-Bird AMU	Enforced WAN: WAN5 - cbird	Any	IP Network 192.168.100.0/24	Any	✕
hsn-router-access-lor	Enforced WAN: WAN3 - hsn	Any	IP Network 210.5.235.24/29	Any	✕
hsn-router-access	Enforced WAN: WAN3 - hsn	Any	IP Network 137.110.254.216/29	Any	✕
ucsd-voip-traffic-must-...	Priority WAN: WAN1 - sho...	IP Address 137.110.149.8	Any	Any	✕
sophos-update-traffic-...	Enforced VPN: revelle	IP Address 137.110.149.196	Any	TCP Any	✕
boat-traffic-must-origin...	Priority VPN: revelle...	Any	IP Address 137.110.151.12	Any	✕
terminate-persistent-lo...	Priority WAN: WAN1 - sho...	IP Address 172.19.149.1	Any	TCP Any	✕
ssh	Priority VPN: revelle...	Any	Any	TCP 22	✕
smtp	Priority VPN: revelle...	Any	Any	TCP 25	✕
pop	Overflow WAN1 - shore cab...	Any	Any	TCP 110	✕
imap	Overflow WAN1 - shore cab...	Any	Any	TCP 143	✕
smtps	Priority VPN: revelle...	Any	Any	TCP 465	✕
imaps	Priority VPN: revelle...	Any	Any	TCP 993	✕
pops	Overflow WAN1 - shore cab...	Any	Any	TCP 995	✕
Default	Priority WAN: WAN1 - shore cable...				

Add Rule

smtp

Service Name *: smtp
Enable: ☑
Source: Any
Destination: Any
Protocol: TCP ⇕ ← ☐ Protocol Selection Tool
Port *: Single Port ⇕ Port: 25
Algorithm: Priority
Priority Order: Highest Priority — VPN: revelle; WAN: WAN1 - shore cable; WAN: WAN2 - wifi-as-wan; WAN: Mobile Internet; WAN: WAN3 - hsn; WAN: WAN4 - fbb; WAN: WAN5 — Lowest Priority. Not In Use
Terminate Sessions on Link Recovery: Enable

Save Cancel

This rule means traffic destined off-ship for port 25 (smtp) should go first via the shore cable, if that is available, and if somehow its not (maybe UCSD campus lost power) then it should go via the shore cable . . . if that's not available then via wifi-as-wan — and so on . .

Algorithm: Priority
Priority Order
Terminate Sessions on Link Recovery

Help Close

This field allows you to choose the WAN selection algorithm.

Weighted Balance - Traffic will be proportionally distributed among available WAN connections according to the specified load distribution weight.

Persistence - Traffic coming from the same machine will be persistently routed through the same WAN connection.

Enforced - Traffic will be routed through the specified WAN connection regardless of the connection's health status.

Priority - Traffic will be routed through the healthy WAN connection that has the highest priority.

Overflow - Traffic will be routed through the healthy WAN connection that has the highest priority until it is in full load. When this connection gets saturated, new sessions will be routed to the next healthy WAN connection that is not in full load.

Least Used - Traffic will be routed through the healthy WAN connection that is selected in the field *Connection* and has the most available downlink bandwidth.

Lowest Latency - Traffic will be routed through the healthy WAN connection that is selected in the field *Connection* and has the lowest latency. Periodic latency checking packets are sent to the selected connections to determine their latency values. Thus additional network usage will be incurred.

Different algorithms are available for deciding which WAN to use for outbound traffic — and can be customized for different traffic types easily via the web-ui

Built-in and accessible monitoring/diagnostics?

Yes via web interface (and ssh) — very useful.

WAN 1 — pier-lan
IP Address: 137.110.255.77 Details... Status: Connected Disconnect

WAN 2 — wifi-as-wan
IP Address: 192.168.20.10 Details... Status: Connected Disconnect

WAN 3 — fbb-lan
IP Address: 216.86.248.54 Details... Status: Connected Disconnect

Mobile Internet
IP Address: 166.140.38.135 Details... Status: Connected to Verizon Wireless

LAN Interface
Router IP Address: 172.19.148.254

SpeedFusion™ Status
Group: Established

Device Information

Model:	Peplink Balance 380
Firmware:	5.4.9m24 build 2655
Uptime:	63 days 11 hours 2 minutes
CPU Load:	85%
Throughput:	⇓ 0.0 Mbps ⇑ 0.0 Mbps

Copyright © Peplink. All rights reserved.

Client List ☑ Online Clients Only

IP Address	Name	Download (kbps)	Upload (kbps)	MAC Address	Import
137.110.148.170		0	0		
137.110.149.1		0	0		
137.110.149.4		0	0		
137.110.149.5		0	0		
137.110.149.8		0	0		
137.110.149.12		0	0		
137.110.149.25		0	0		
137.110.149.26		0	0		
137.110.149.31		18	26		
137.110.149.32		0	0		
137.110.149.33		8	0		
137.110.149.34		0	0		
137.110.149.35		0	0		
137.110.149.36		0	0		
137.110.149.37		0	0		
137.110.149.38		0	0		
137.110.149.39		16	3		
137.110.149.40		0	0		
137.110.149.41		0	0		
137.110.149.42		0	0		
137.110.149.43		0	0		
137.110.149.44		0	0		
137.110.149.45		0	0		
137.110.149.46		0	0		
137.110.149.47		0	0		
137.110.149.50		0	0		
137.110.149.51		0	0		
137.110.149.52		0	0		
137.110.149.53		0	0		
137.110.149.54		0	0		
137.110.149.55		0	0		
137.110.149.56		0	0		
137.110.149.58		0	0		
137.110.149.59		0	0		
137.110.149.60		0	0		
137.110.149.61		0	0		

WAN3 - hsn [Show Details]

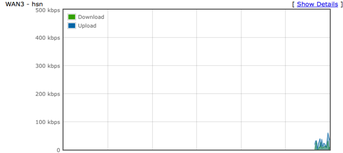

Avg: ⇓ 10.78 kbps ⇑ 26.22 kbps Peak: ⇓ 31 kbps ⇑ 60 kbps Stacked ☐

	Download	Upload	Total
Overall	10 kbps	34 kbps	44 kbps

WAN4 - fbb [Hide Details]

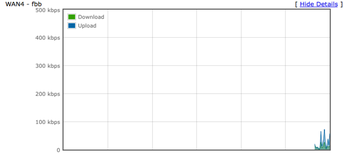

Avg: ⇓ 10.50 kbps ⇑ 23.33 kbps Peak: ⇓ 26 kbps ⇑ 72 kbps Stacked ☐

	Download	Upload	Total
Overall	13 kbps	57 kbps	70 kbps
HTTP	0 kbps	0 kbps	0 kbps
HTTPS	0 kbps	0 kbps	0 kbps
IMAP	0 kbps	0 kbps	0 kbps
POP3	0 kbps	0 kbps	0 kbps
SMTP	0 kbps	0 kbps	0 kbps
Others	13 kbps	57 kbps	70 kbps

Mobile Internet [Show Details]

Outbound

Protocol	Source IP	Destination IP	Service	Interface	Idle Time
TCP	192.168.149.91:36429	67.23.37.23:443	HTTPS/SSL	WAN3 - hsn	00:00:01
TCP	192.168.149.91:36426	67.23.37.23:443	HTTPS/SSL	WAN3 - hsn	00:00:12
TCP	192.168.149.91:36427	67.23.37.23:443	HTTPS/SSL	WAN3 - hsn	00:00:12
TCP	172.19.149.1:45041	103.29.28.153:443	HTTPS	WAN3 - hsn	00:00:14
TCP	192.168.149.91:36428	67.23.37.23:443	HTTPS/SSL	WAN3 - hsn	00:00:14
TCP	192.168.149.91:48072	74.125.224.40:443	HTTPS	WAN3 - hsn	00:00:28
TCP	137.110.149.32:49778	74.125.224.39:443	HTTPS	WAN3 - hsn	00:00:34
TCP	137.110.149.32:64036	132.239.50.202:993	Secure IMAP	VPN	00:00:46
TCP	137.110.149.32:64055	132.239.50.202:993	Secure IMAP	VPN	00:00:49
TCP	137.110.149.32:55326	17.173.66.181:443	HTTPS	WAN3 - hsn	00:00:51
TCP	137.110.149.32:60150	74.125.224.36:443	HTTPS	WAN3 - hsn	00:00:53
TCP	137.110.149.32:64054	132.239.50.202:993	Secure IMAP	VPN	00:00:54
TCP	137.110.149.32:64049	132.239.50.202:993	Secure IMAP	VPN	00:00:57
TCP	192.168.149.91:37511	74.125.224.46:443	HTTPS	WAN3 - hsn	00:00:57
TCP	137.110.149.32:62057	LAN Client		WAN3 - hsn	00:00:59
TCP	192.168.149.91:50748	74.125.239.148:443	HTTPS	WAN3 - hsn	00:00:59

Total searched results: 16

Inbound

Protocol	Source IP	Destination IP	Service	Interface	Idle Time
			No sessions		

Total searched results: 0

Transit

Protocol	Source IP	Destination IP	Service	Interface	Idle Time
			No sessions		

Total searched results: 0

Every user is assigned to a group. A member of ‘group A’ will be given a quota of X megabytes per day, a member of ‘group B’ gets Z megabytes per day, and so on

Cyberoam is a full featured firewall/security appliance with a slew of features including powerful deep packet inspection capabilities. It uses a captive portal technique to discover ‘user-identity’ — ‘user-identity’ can then be used as part of any traffic identification rule

This allows us to define -

- per user per day bandwidth quotas
- firewall rules, quotas, and any other policy can take into account the identity of the user at the other end — we don’t care about tracking devices

This is much easier to manage and allows specification of better, more equitable as well as more adaptable wan access control policies.

The cyberoam provided documentation is not good — but the implementation is straightforward once you understand how its UI is organized.

We moved all of the logic to control ‘who’ has access to ‘what’ via the ship’s WANs into a single security appliance called cyberoam (approx $500 hardware purchase - plus $100 per year for subscription to url filtering and layer-7 identitication/ categorization lists)

What is a captive portal?

html login page + state which tracks whether or not a user is logged in + login state-aware http redirector

Login is required to access the internet from Revelle.

Please NOTE -- the *Login* button changes to a *Logout* button after you have successfully authenticated.

You may only be logged in from one device at a time -- if you need to login from a second device first close the internet session on the first device by logging out via the LOGOUT button on this page.

KEEP this browser page OPEN and return here to LOGOUT when you are done using the internet. If you close this page and wish to return to it -- use the link posted on the ship webpage -- http://ship

Revelle Internet Access Management

Username

admin

Password

••••••••••••

Login

- Click here for account details

Use mobile sites rather than their desktop counterparts!

m.facebook.com

m.gmail.com

m.yahoo.com

twitter

amazon

Mobile apps are also generally better to use than websites designed for the desktop.

By logging into this page you acknowledge that you understand that network connectivity from the ship is via expensive, slow and often congested satellite internet. You accept that this system is installed on this ship for business-critical scientific and shipboard communications -- all other use is ancillary and subject to being blocked at any time. You agree to take care with your usage of the link, and to take every action available to you to limit your usage so that the system is likely to perform its critical functions when required. You also certify that all your devices comply with ucsd minimum security standards.

Cyberoam implementation overview

1. Define a set of objects to identify subnets and any hosts or ip segments which might need special treatment

Objects –Hosts –ip hosts

	Name	Type	Value	Detailed Description	Zone	Host Policy
1	DC1_VIP	Network	Net		Any	Allow
2	DC2_VIP	Network	Net		Any	Allow
3	DC1_VIP_TCP	Network	Net		Any	Allow
4	DC1-01	Network	Net	172.16.0.0/255.255.255.255	Any	Allow
5	DC1-02	Network	Net	172.16.0.0/255.255.255.255	Any	Allow
6	DC1-03	Network	Net		Any	Allow
7	DC1-04	Network	Net		Any	Allow
8	10.Net-internal	Network	10.0.0.0/255.255.255.0		Any	Allow
9	10.Net-internal-2	Network	10.0.0.0/255.255.255.0		Any	Allow
10	10.Net-public	Network	10.0.0.0/255.255.255.0		Any	Allow
11	10.Net-DMZ	Network	10.0.0.0/255.255.255.0		Any	Allow
12	172.16.Net-internal	Network	172.16.0.0/255.255.255.0		Any	Allow
13	172.16.Net-public	Network	172.16.0.0/255.255.255.0		Any	Allow
14	172.16.Net-DMZ	Network	172.16.0.0/255.255.255.0		Any	Allow
15	172.16.Net-Vmware	Network	172.16.0.0/255.255.255.0		Any	Allow
16	172.16.Net-Vmware-2	Network	172.16.0.0/255.255.255.0		Any	Allow
17	172.16.Net-Vmware-3	Network	172.16.0.0/255.255.255.0		Any	Allow
18	172.16.Net-Vmware-4	Network	172.16.0.0/255.255.255.0		Any	Allow

and for the general firewall rules, define a set of objects representing well known service ports for common traffic types

Objects –Services –Services

	Name	Type	Details	Detailed Description	Protocol
1	CIFS/TCP	TCP/UDP	TCP:137,139,445		CIFS
2	CIFS/UDP	TCP/UDP	UDP:137,138,139,445		CIFS
3	DHCP-UDP	UDP	UDP:67, 68	DHCP, 67 - Server, 68 - DHCP Client	
4	DNS	TCP/UDP	TCP:53,UDP:53		DNS
5	FTP	TCP/UDP	TCP:21,UDP:21	FTP, TCP:21 - Control, UDP:21 - Control	FTP
6	HTTP	TCP/UDP	TCP:80,UDP:80	HTTP, TCP:80 - Data, UDP:80 - Data	HTTP
7	HTTPS-TCP/UDP	TCP/UDP	TCP:443,UDP:443		
8	IMAP-TCP/UDP	TCP/UDP	TCP:143,UDP:143	IMAP, TCP:143 - Control, UDP:143 - Control	
9	IMAPS-TCP/UDP	TCP/UDP	TCP:993,UDP:993		
10	LDAP	TCP/UDP	TCP:389,UDP:389		
11	LDAPS	TCP/UDP	TCP:636,UDP:636		
12	MS-SQL	TCP/UDP	TCP:1433,UDP:1434	MS-SQL, TCP:1433 - Data, UDP:1434 - Lookup	
13	NFS	TCP/UDP	TCP:2049,UDP:2049		
14	NTP	TCP/UDP	TCP:123,UDP:123		
15	POP3	TCP/UDP	TCP:110,UDP:110		
16	POP3S	TCP/UDP	TCP:995,UDP:995		
17	RDP	TCP/UDP	TCP:3389,UDP:3389		
18	RPC	TCP/UDP	TCP:135,UDP:135		
19	SMB	TCP/UDP	TCP:445,UDP:445		
20	SMTP	TCP/UDP	TCP:25,UDP:25		
21	SMTPS	TCP/UDP	TCP:465,UDP:465		
22	SNMP	TCP/UDP	UDP:161,UDP:162		
23	SSH	TCP/UDP	TCP:22,UDP:22		
24	TELNET	TCP/UDP	TCP:23,UDP:23		
25	TFTP	TCP/UDP	UDP:69,UDP:1024-1050		
26	VNC	TCP/UDP	TCP:5900,UDP:5900		
27	WINS	TCP/UDP	TCP:1512,UDP:1512		
28	XDMCP	TCP/UDP	UDP:177		

and also define some other objects like the quotas you will need:

Identity –> Policy –> Data transfer

	Name	Source Type	Source Value	Destination Type	Destination Value	Protocol	Action	Quota	Schedule	Log
1	Quota-DC1-admin-backup	Network	10.0.0.0/255.255.255.0	Network	172.16.0.0/255.255.255.0	TCP	Allow	10	Always	Off
2	Quota-DMZ-backup	Network	10.0.0.0/255.255.255.0	Network	172.16.0.0/255.255.255.0	TCP	Allow	10	Always	Off
3	Quota-DC1-backup-public	Network	10.0.0.0/255.255.255.0	Network	172.16.0.0/255.255.255.0	TCP	Allow	10	Always	Off
4	Quota-DC1-backup-DMZ	Network	10.0.0.0/255.255.255.0	Network	172.16.0.0/255.255.255.0	TCP	Allow	10	Always	Off
5	Quota-DC-backup	Network	10.0.0.0/255.255.255.0	Network	172.16.0.0/255.255.255.0	TCP	Allow	10	Always	Off

2. organize the ip objects in your network into semantically meaningful groups

Objects –Hosts –ip Host Group

	Group Name	Group Members	Description	Host Policy
1	10.Net-public	10.Net-public, 10.Net-public-2, 10.Net-public-3, 10.Net-public-4, 10.Net-public-5, 10.Net-public-6, 10.Net-public-7, 10.Net-public-8, 10.Net-public-9, 10.Net-public-10		Allow
2	10.Net-public-subnets	10.Net-public-2, 10.Net-public-3, 10.Net-public-4, 10.Net-public-5, 10.Net-public-6, 10.Net-public-7, 10.Net-public-8, 10.Net-public-9, 10.Net-public-10		Allow
3	10.Net-public-subnets-2	10.Net-public-3, 10.Net-public-4, 10.Net-public-5, 10.Net-public-6, 10.Net-public-7, 10.Net-public-8, 10.Net-public-9, 10.Net-public-10		Allow
4	172.16.Net-DMZ	172.16.Net-DMZ-Vmware, 172.16.Net-DMZ-Vmware-2, 172.16.Net-DMZ-Vmware-3, 172.16.Net-DMZ-Vmware-4		Allow
5	172.16.Net-DMZ-Vmware	172.16.Net-DMZ-Vmware-2, 172.16.Net-DMZ-Vmware-3, 172.16.Net-DMZ-Vmware-4		Allow
6	172.16.Net-DMZ-Vmware-2	172.16.Net-DMZ-Vmware-3, 172.16.Net-DMZ-Vmware-4		Allow
7	172.16.Net-DMZ-Vmware-3	172.16.Net-DMZ-Vmware-4		Allow
8	172.16.Net-DMZ-Vmware-4			Allow

and organize the services into service groups:

Objects –Services –Services Group

	Name	Description	Member
1	Web-access	TCP:443,80,8080, UDP:443,80,8080	HTTP, HTTPS-TCP/UDP, Proxy-8080
2	Mail-access-all-protocols	all mail protocols protocols to allow or not	IMAP-TCP/UDP, IMAPS-TCP/UDP, POP3, POP3S, SMTP, SMTPS
3	Mail-access-secure-only	secure mail protocols protocols to allow or not	IMAPS-TCP/UDP, POP3S, SMTPS
4	Mail-access-secure-only-2	secure mail protocols protocols to allow or not	IMAPS-TCP/UDP, POP3S, SMTPS
5	Windows-services	Traffic which should be allowed for windows to work	CIFS/TCP, CIFS/UDP, DHCP-UDP, DNS, NTP
6	Windows-services-2	Traffic which should be allowed for windows to work	CIFS/TCP, CIFS/UDP, DHCP-UDP, DNS, NTP
7	Web-admin	Web admin protocols and ports for firewall	HTTP, HTTPS-TCP/UDP, SSH, TELNET
8	Mail-admin	mail administration	POP3, POP3S, SMTP, SMTPS

3. Setup user groups – figure out how many different groups of users you have, and define how each group should be treated by the network

Objects –Identity –Groups

	Group Name	Group Type	Description	Lan Policy	Wan Policy
1	Admin-Group	Security Group		Allow	Allow
2	Guest	Security Group		Allow	Allow
3	Lab-1	Security Group		Allow	Allow
4	Lab-2	Security Group		Allow	Allow
5	Contractors	Security Group		Allow	Allow
6	Sales	Security Group		Allow	Allow
7	HR	Security Group		Allow	Allow
8	Marketing	Security Group		Allow	Allow
9	Internet_Unlimited	Security Group		Allow	Allow
10	Internet_Quota	Security Group		Allow	Allow
11	Internet_Time	Security Group		Allow	Allow
12	Internet_Quota_Time	Security Group		Allow	Allow
13	Internet_Restricted	Security Group		Allow	Allow
14	Internet_Restricted_Quota	Security Group		Allow	Allow
15	Internet_Restricted_Time	Security Group		Allow	Allow
16	Internet_Restricted_Quota_Time	Security Group		Allow	Allow

Group edit ur

Group Name: Admin-Group
Group Type: Security Group

Members

	Name	Type
1	Admin-Group	Security Group
2	Application-Admin	general user Filter
3	Surfing-Admin	Internet Access Policy
4	Web-Filter-Admin	Web Filter Policy
5	Mail-Admin	Mail Filter Policy
6	Data-Admin	Data Transfer Policy
7	Ftp-Admin	Ftp Access Policy
8	Vpn-Admin	Vpn Access Policy
9	Wifi-Admin	Wireless Access Policy
10	Firewall-Admin	Firewall Policy
11	Admin-Group-Vmware	Active Directory Group
12	Admin-Group-Vmware-Group	Active Directory Group

OK | Add Members | Show Group Members | Cancel

4. Define your inbound firewall – this is to control what traffic from outside of your network cyberoam will allow to initiate connections to inside of your network

Reference the ongoing defined series to define firewall rules that are straightforward and clear. Make rules as narrow and specific as possible. For ease of debugging avoid creating more than one rule with the same source/destination ip address

Firewall –> Rule –> WAN-LAN

	Name	Source	Dest	Service	Action	Schedule	Description	Log	Status
WAN - LAN - Total: 11									
1	c-traffic-mail-secure-only-2	Any	Any	Mail-access-secure-only-2	Accept	Always		Off	On
2	c-traffic-mail-secure-only	Any	Any	Mail-access-secure-only	Accept	Always		Off	On
3	c-traffic-mail-all-protocols	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
4	c-traffic-mail-all-protocols-2	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
5	c-traffic-Internet-access	Any	Any	Web-access	Accept	Always		Off	On
6	c-traffic-Internet-access-2	Any	Any	Web-access	Accept	Always		Off	On
7	c-traffic-Internet-access-3	Any	Any	Web-access	Accept	Always		Off	On
8	c-traffic-Internet-access-4	Any	Any	Web-access	Accept	Always		Off	On
9	c-traffic-Internet-access-5	Any	Any	Web-access	Accept	Always		Off	On
10	c-traffic-Internet-access-6	Any	Any	Web-access	Accept	Always		Off	On
11	c-traffic-Internet-access-7	Any	Any	Web-access	Accept	Always		Off	On

Firewall –> Rule –> WAN-LAN

	Name	Source	Dest	Service	Action	Schedule	Description	Log	Status
WAN - LAN - Total: 11									
1	c-traffic-mail-secure-only-2	Any	Any	Mail-access-secure-only-2	Accept	Always		Off	On
2	c-traffic-mail-secure-only	Any	Any	Mail-access-secure-only	Accept	Always		Off	On
3	c-traffic-mail-all-protocols	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
4	c-traffic-mail-all-protocols-2	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
5	c-traffic-Internet-access	Any	Any	Web-access	Accept	Always		Off	On
6	c-traffic-Internet-access-2	Any	Any	Web-access	Accept	Always		Off	On
7	c-traffic-Internet-access-3	Any	Any	Web-access	Accept	Always		Off	On
8	c-traffic-Internet-access-4	Any	Any	Web-access	Accept	Always		Off	On
9	c-traffic-Internet-access-5	Any	Any	Web-access	Accept	Always		Off	On
10	c-traffic-Internet-access-6	Any	Any	Web-access	Accept	Always		Off	On
11	c-traffic-Internet-access-7	Any	Any	Web-access	Accept	Always		Off	On

You take an active hit firewall rule BEACT actions to block traffic from specific ip addresses that are trying to access your network. Cyberoam has the ability to automatically block traffic based on rules you define. This is a very powerful feature that can be very useful in administering the network security.

The active hits and counts are the amount of time traffic is matched against the firewall rule. Active hits and counts are compared rule after or another to correlation failed between the firewall entries list.

5. Define your outbound firewall – this is to control what traffic from outside of your network cyberoam will allow to initiate connections to inside of your network

Firewall –> Rule –> LAN-WAN

	Name	Source	Dest	Service	Action	Schedule	Description	Log	Status
LAN - WAN - Total: 11									
1	c-traffic-mail-secure-only-2	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
2	c-traffic-mail-secure-only	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
3	c-traffic-mail-all-protocols	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
4	c-traffic-mail-all-protocols-2	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
5	c-traffic-Internet-access	Any	Any	Web-access	Accept	Always		Off	On
6	c-traffic-Internet-access-2	Any	Any	Web-access	Accept	Always		Off	On
7	c-traffic-Internet-access-3	Any	Any	Web-access	Accept	Always		Off	On
8	c-traffic-Internet-access-4	Any	Any	Web-access	Accept	Always		Off	On
9	c-traffic-Internet-access-5	Any	Any	Web-access	Accept	Always		Off	On
10	c-traffic-Internet-access-6	Any	Any	Web-access	Accept	Always		Off	On
11	c-traffic-Internet-access-7	Any	Any	Web-access	Accept	Always		Off	On
12	c-traffic-mail-secure-only-2	Any	Any	Mail-access-secure-only-2	Accept	Always		Off	On
13	c-traffic-mail-secure-only	Any	Any	Mail-access-secure-only	Accept	Always		Off	On
14	c-traffic-mail-all-protocols	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
15	c-traffic-mail-all-protocols-2	Any	Any	c-smtp-proxy(80)	Drop	Always		Off	On
16	c-traffic-Internet-access	Any	Any	Web-access	Accept	Always		Off	On
17	c-traffic-Internet-access-2	Any	Any	Web-access	Accept	Always		Off	On
18	c-traffic-Internet-access-3	Any	Any	Web-access	Accept	Always		Off	On
19	c-traffic-Internet-access-4	Any	Any	Web-access	Accept	Always		Off	On
20	c-traffic-Internet-access-5	Any	Any	Web-access	Accept	Always		Off	On
21	c-traffic-Internet-access-6	Any	Any	Web-access	Accept	Always		Off	On
22	c-traffic-Internet-access-7	Any	Any	Web-access	Accept	Always		Off	On

Firewall rule matching occurs from top to bottom as with a traditional firewall rule – the rule associated with the first matching pattern is applied to the traffic.

Unlike a traditional firewall rule "user identity" can be used as part of the match criteria by checking the "user identity" box and selecting a list of groups or users to match.

Traffic from an unauthenticated source will have match any rules which have "Match Identity" checked.

General Settings

Rule Name: c-traffic-Internet-access-7

Rule Description: This rule allows all traffic from the internal network to the Internet

Rule Settings

Status: On
Action: Accept
Source Zone(s): LAN
Source Network: All-networks-on-internal-LAN, All-Internal-public-networks, Vmware-1, Vmware-2, Vmware-3, Vmware-4
Destination Zone(s): Any
Destination Network: Any
Service: Web-access, Proxy-8080
Schedule: Always
Log Traffic: Off

Advanced Settings Security Policy Match Criteria, Log Traffic

Security Policy Match Criteria

Application Filter: Dont-Calling-applicat
Web Filter: Dont-Calling-applicat
IPS: None
Traffic Shaping: None
QoS: Qos-Internet
Qos in Scheduling: Qos1, Qos2, Qos3, Qos4, Qos5, Qos6, Qos7, Qos8

Match Identity: Match Identity
User or Group Name: Internet-Web-Filtering
Time Through Gateway: Disabled
Remote Gateway: Any

Log Traffic

Log Firewall Traffic: Off
Log Illegal Traffic: Off

OK | Cancel

Match Identity

Match Identity: Match Identity

If an unauthenticated http request goes through the whole chain without being accepted, dropped, or rejected – it is then rejected with a redirect to the Cyberoam "Login" page. This is the "Internet Access Policy".

Cyberoam guest page access to non-authenticated user and http redirect to access page of the http requests.

Login to request to access the Internet Pass Results

Please NOTE – the Login button changes to a Logout button after you have successfully authenticated.

You may not be logged in Cyberoam Internet Access Policy. If you think that you should have access to Internet then click the button below to login to the Cyberoam Internet Access Policy.

VERY IMPORTANT! Click the button to LOGOUT when you are done using the Internet. If you close the page and click back or exit it – you may be considered to be still logged in.

Click here to login

Novell Internet Access Management

Username:
Password:
Domain:
Login
Cancel

Enter Novell Net account details

You can control how little login page looks

System –> Configuration –> Captive Portal

General Settings

Title: Cyberoam Internet Access Policy
Welcome Text: Click the Login button to access the Internet
Logo URL: http://www.cyberoam.com
Page Title: Bad Behavior Internet Access Management
Logo Web Page Hyperlink: Enter a URL in the format http://www.domain.com or https://www.domain.com or ftp://ftp.domain.com
Logo Page Hyperlink: Enter a URL in the format http://www.domain.com or https://www.domain.com or ftp://ftp.domain.com
Information Submit: Disabled
Disclaimer Required: Disabled
Legal Notice Caption: Legal
Legal Notice Text: Legal
Title Text Information: Enter text for special details

Color Scheme

Background Color: 000000
Header Center Text Color: 000000
Text Color: 000000
Text Hyperlink Color: 000000
Header Center Text Color: 000000
Text Hyperlink Color: 000000

Apply | Preview >> | Reset To Default

Use mobile sites rather than their desktop counterparts!

m.facebook.com
m.gmail.com
m.yahoo.com
twitter
twitter
twitter

Mobile apps are also generally better to use than websites designed for the desktop.

The Mobile Web Pages are accessibility that is considered the standard recommendation for the future. For instance, you can click a button to switch to a version of the mobile site optimized for accessibility. Mobile Web Pages provides a form of language or content negotiation. The client sends a header to the server indicating what it can accept. The server then chooses a version of the resource based on the client's capabilities. Mobile Web Pages provides a form of content negotiation. Mobile Web Pages provides a form of content negotiation. Mobile Web Pages provides a form of content negotiation