

LOAD BALANCING WITH EXCHANGE 2010

David Zazzo
Senior Consultant
Microsoft Corporation



Agenda


- History of Load Balancing
- Load Balancing Principles
- Exchange 2010 Load Balancing Requirements
- Hardware Load Balancing Vendors

History of Load Balancing

- **Software Load Balancing**
 - Windows NT 4.0 – Windows Network Load Balancing services introduced
- **Issues with WNLB**
 - Switch/Port flooding in Unicast mode
 - NAT / Source IP pool
 - Scalability over 8 nodes
 - Service awareness
 - Not supported with Windows Failover Clustering
 - Add/remove single node causes all clients to reconnect

Load Balancing Principles

- Dedicated hardware load balancers



- Load Balancer -> Application Delivery Controller (ADC)
 - Application health, granular persistence methods

Load Balancing Principles

- **Application Delivery Controller Benefits**
 - Load distribution based on target server capacity / connection information
 - “Smart” compression
 - Don't compress for low latency/high bandwidth clients
 - SSL Offload / Bridging
 - Caching of OWA attachments / items
 - Rate shaping and TCP packet stream modifications
 - Dynamic window sizing on a per client/per server basis

Load Balancing Principles

- **Routing Options**
 - SNAT
 - Load Balancer Default Gateway
 - Direct Server Return
- **Hardware load balancing terminology**
 - One-Armed
 - Two-Armed
 - Node / Host / Member / Server
 - Persistence
 - DNAT

Load Balancing Principles

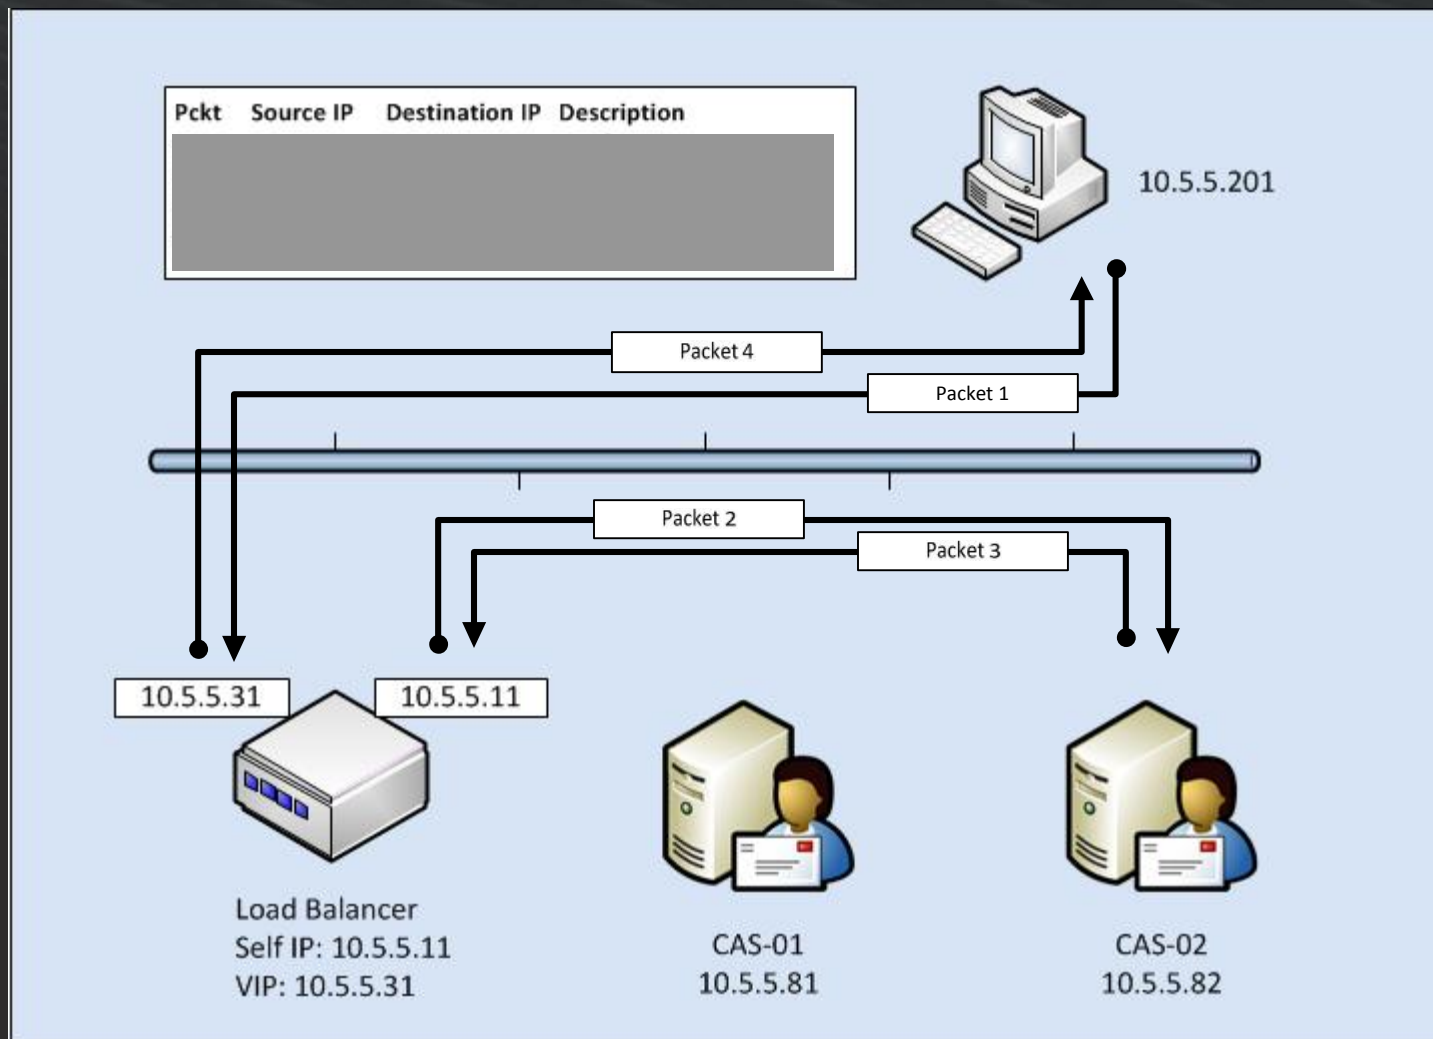
- **One Armed**

- Single network, VIP is same subnet as clients and servers
- Hardware Load Balancer is an independent node
 - May have multiple IPs (management interface, etc...)
- Single VLAN
- Requires SNAT or server will respond directly to client
 - This would break the security as client is not expecting a return from the server IP but from the VIP

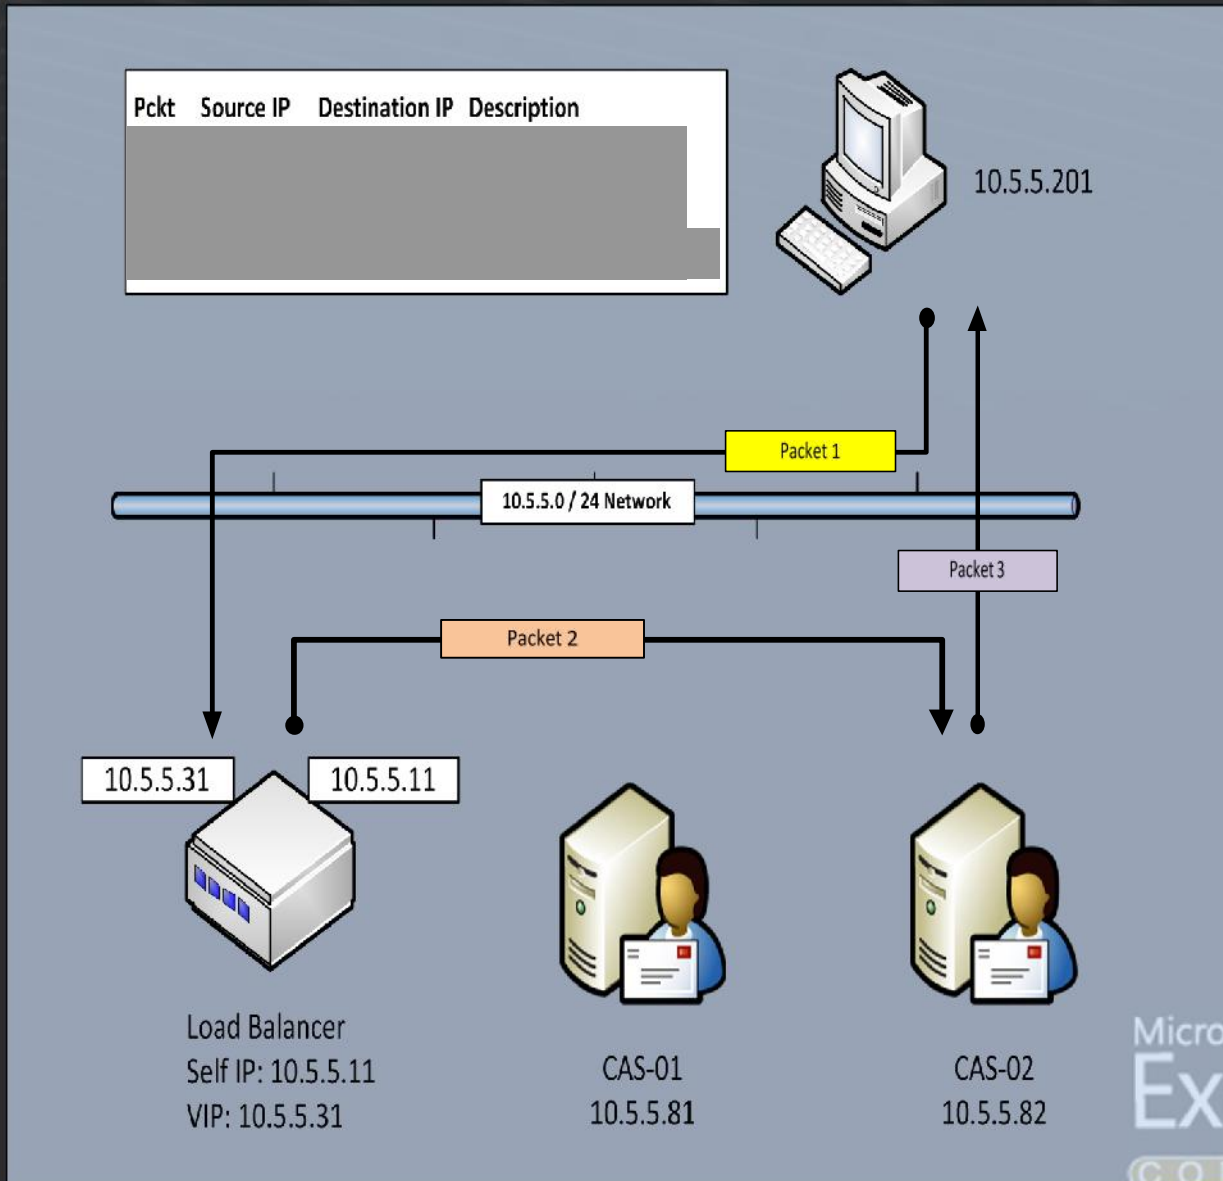
- **Two Armed**

- 2 Networks / VLANS
- Dedicated VLAN / network for Load Balanced servers / array
- Load Balancer “straddles”
- Can use SNAT or Load Balancer Default Gateway

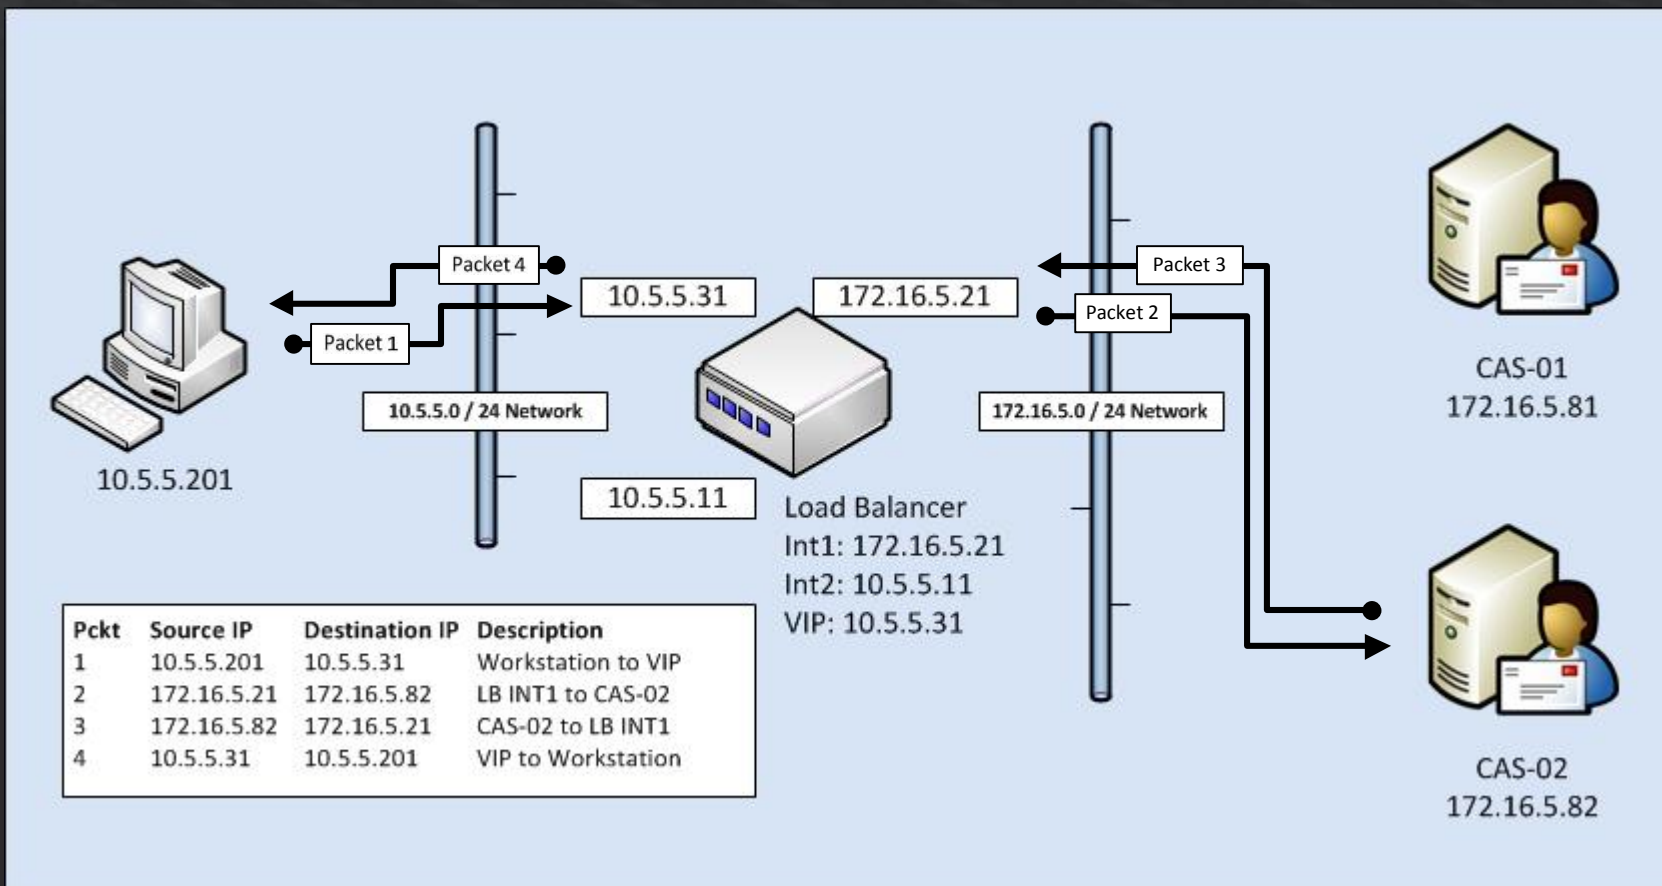
Single Arm - SNAT



Single Arm - Direct Server Return



Dual Arm – SNAT



Direct Server Return

- DSR leverages MAT (MAC Address Translation) and modified nodes
- Server responds with the LB VIP address to the clients
- DSR Requires
 - Configure the virtual service on load balancer to route client packets (instead of NAT'ing them)
 - Configure a loopback adapter on each CAS server
 - Add the Load Balancer VIP to the loopback adapter
 - Use Layer 4 (Layer 7 features not supported – persistence, caching, compression, etc...)
- Requires additional complexity, easier to use SNAT instead

Load Balancing Principles

- SNAT versus SNAT “Pool”
- SNAT (1 IP) is limited to 65,535 concurrent connections to each target server
 - NAT maps source IP/source port to target IP/target port
- SNAT Pools (2+ IPs) are limited to 65,535 concurrent connections per source SNAT IP to each target server
 - Some Hardware Load Balancers allow SNAT “pools” to allow additional concurrent connections

Load Balancing Principles

- MAPI RPC (RPC CAS) requires the same source IP address for all RPC connections from that client
 - Load Balancer MUST use same SNAT source IP from the pool for each client
- MAPI RPC Client Request 1 -> SNAT IP1
- MAPI RPC Client Request 2 -> SNAT IP2 = Broken
 - Fix 1: Turn off SNAT Pool and use single SNAT IP
 - or
 - Fix 2: Configure LB to maintain client/SNAT IP when using MAPI RPC

Load Balancing Principles

- **Node / Host / Member / Server**
 - Node / Host / Server represent an IP address and/or physical server
 - Member is an IP address + port combination
 - IP address + TCP 443 for example
- **Client Logging**
 - SNAT hides client IP address from server logs

Summary of Routing Options

SNAT

- No modifications to CAS
- Return traffic passes through LB
- Most common
- Client IP not on server logs

DSR

- Requires modifying CAS servers
- Return traffic doesn't pass through LB
- Reduces persistence options
- Introduces timeout complexity
- Rarely used (typically large media hosters)

LBDG

- Requires modifying CAS servers
- Return traffic passes through LB
- Adds complexity to network routing (need static routes)
- Client IP is on server logs

Recommended

Load Balancing Principles

- Persistence / Affinity
 - SSL Session ID
 - Cookie
 - Load Balancer generated “Cookie”
 - Application generated “Cookie”
 - Hybrid
 - Source IP address
 - Universal Persistence
 - Destination Address
 - RDP (not used for Exchange)
 - SIP (not used for Exchange)
- Persistence does not necessarily equal Load Distribution

“Cookie” Persistence (aka the hungry slide)

- All Cookie Persistence methods require SSL termination at LB from the client
- Cannot see persistence records in state table when using cookie LB
 - Insert Cookie
 - Rewrite Cookie
 - Passive Cookie
 - Cookie Hash

Agenda

- History of Load Balancing
- Load Balancing Principles
- Exchange 2010 Load Balancing Requirements
- Hardware Load Balancing Vendors

Load Balancing Traffic Patterns

- **3 types of traffic**

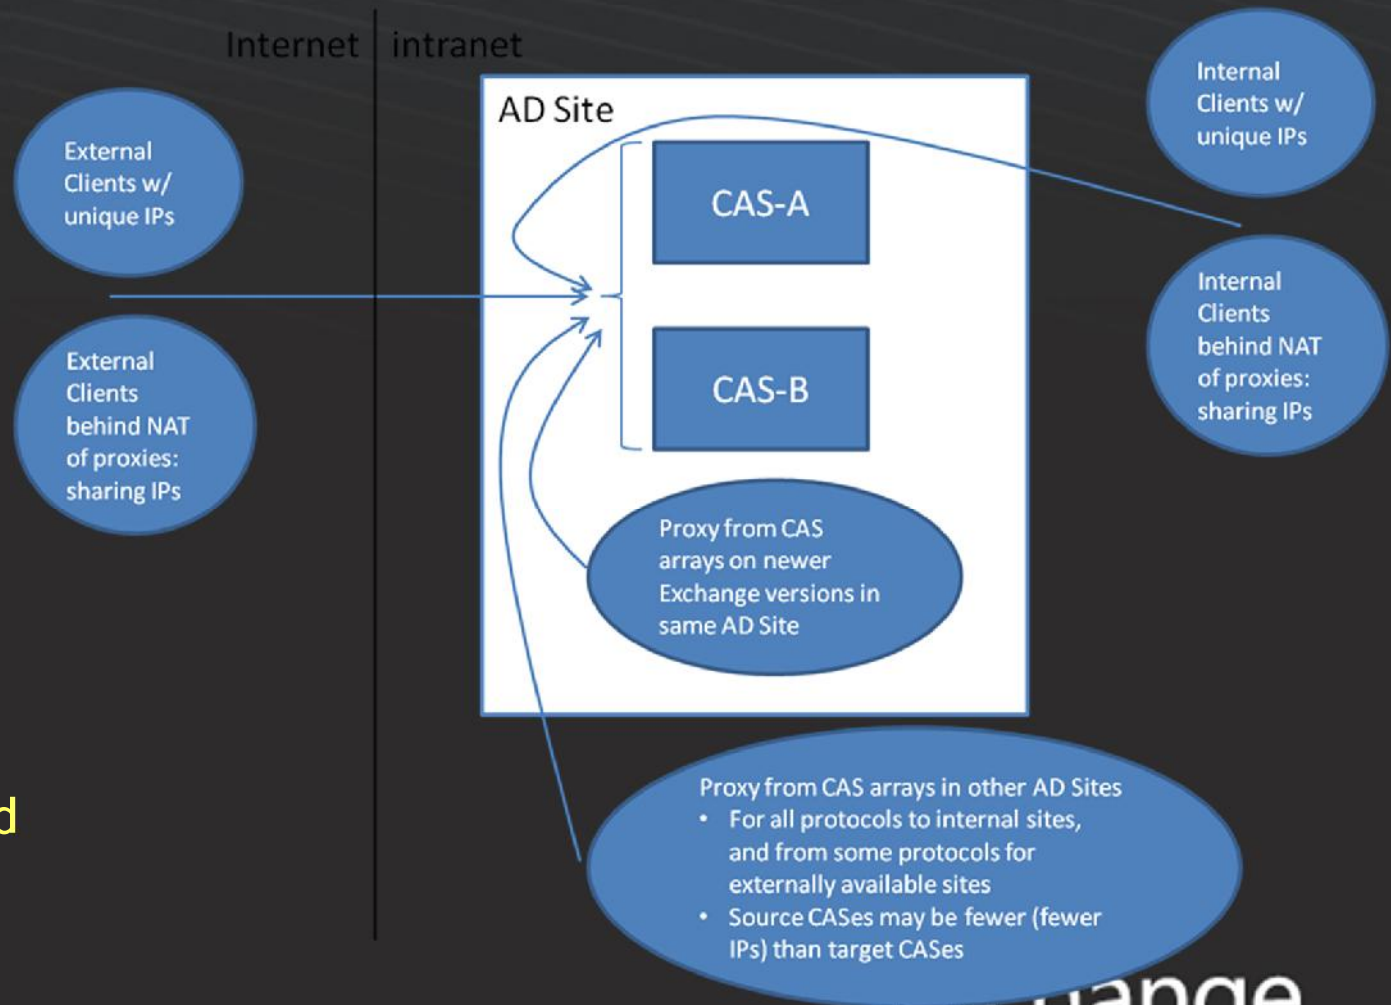
- External
- CAS proxy
- Internal

- **2 entry points**

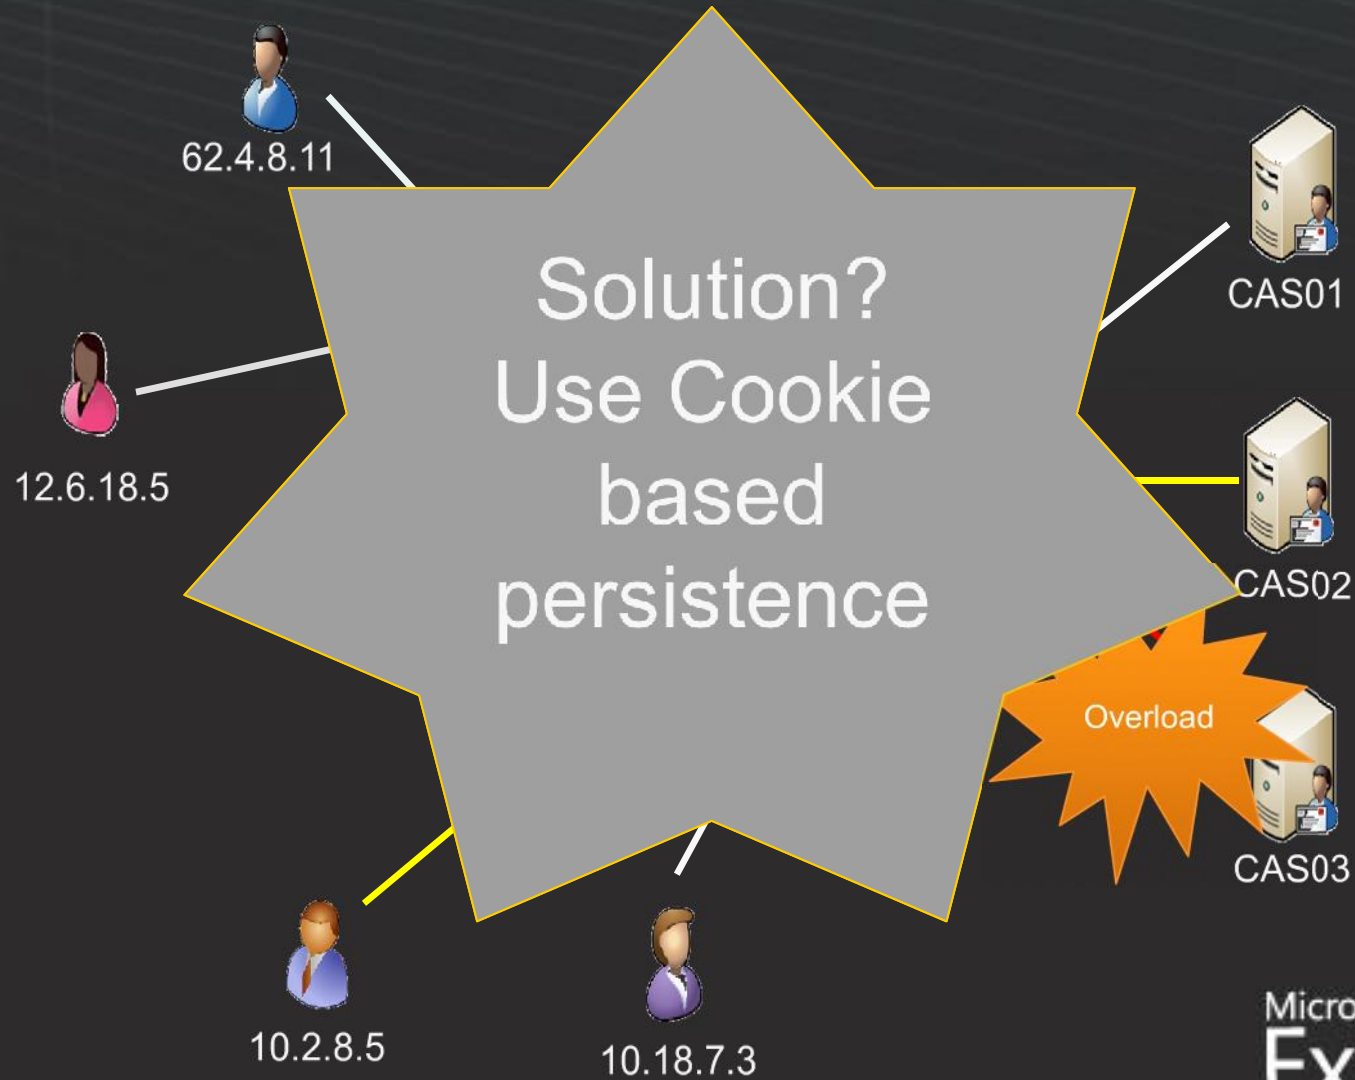
- NAT clients
- non-NAT clients

- **Impacts SSL termination**

- **Impacts FBA and authentication**



NAT and Load Balancing with Source IP



Exchange 2010 Persistence Issues

- **Multiple workload types**
 - OWA / ECP
 - EWS
 - ActiveSync
 - POP/IMAP
 - Outlook Anywhere
 - Outlook MAPI
 - Address Book Service
- **Each has a different client and differing support for persistence methods**
- **Single VIP for all workloads?**
 - Multiple VIPs = easier troubleshooting/logging but requires additional IPs and namespaces

Load Balancing Protocols and Persistence / Affinity / Stickiness

Persistence: Required	Persistence: Recommended	Persistence: Not Required
Outlook Web App	Outlook Anywhere	Offline Address Book
Exchange Control Panel	ActiveSync	AutoDiscover
Exchange Web Services	Address Book Service	POP3
RPC Client Access Service	Remote PowerShell	IMAP4

- **Recommended**
 - Reduced performance without persistence
- **Not required**
 - Does not suffer performance hit without persistence

Exchange 2010 Persistence

OWA and ECP

- OWA and ECP are stateful, require persistence
 - If large source IP range, use source IP persistence
 - Less load on load balancer, no requirement to terminate SSL session
- CAS servers generate a cookie
 - OWA: “UserContext” (“OutlookSession” in SP1)
 - ECP: “msExchEcpCanary”
 - Use application cookie based persistence (requires decrypting the SSL stream)
 - Requires configuration of the LB for each CAS array node
- Load Balancer generated cookie
 - InsertCookie

Exchange 2010 Persistence ActiveSync

- Does not require persistence
 - But... lack of persistence causes server-side performance impact (and increased latency for clients)
- If using Basic Authentication can use the Authorization header
 - Authorization: Basic ZmFrZXVzZXI6eCRwSUFLOUBwOSE=
- ActiveSync Persistence Issues
 - Client IP – all clients may have same IP range from the carrier
 - SSL Session ID – Some ActiveSync clients change the ID often (iPhone)
 - Authorization Header only works with Basic (not client certificate for example)

Exchange 2010 Persistence Exchange Web Services (EWS)

- Stateful protocol, requires persistence
 - Recommend server or load balancer generated cookies
 - May also use SSL session ID
 - If clients don't initiate new SSL session IDs
- EWS Persistence Issues
 - Some clients (OC Phone Edition) don't process cookies

Exchange 2010 Persistence Outlook Anywhere (OA)

- Recommended Persistence Types
 - Client IP
 - “OutlookSession” cookie
 - “OutlookSession” cookie requires OLK 2010
 - No, this isn’t the same cookie as OWA, just named the same
- If persistence isn’t configured for OA
 - CAS Load Balancing Service ensures RPC_DATA_IN and RPC_DATA_OUT align to same CAS
 - Increases performance load on CAS

Exchange 2010 Persistence Offline Address Book (OAB)

- Recommended Persistence Types
 - Client IP
 - SSL Session ID
 - ... or none

Exchange 2010 Persistence

Outlook MAPI / RPC Client Access Service

- **Recommended Persistence Types**
 - Client IP is the only option
 - SSL Session ID, Cookies, etc... are all based on HTTP
- **Use Round Robin Load Balancing method**
- **Static Mapping of RPC ports recommended**
 - Reduces the range of destination ports in load balancer config and memory
 - Easier firewall configuration (if applicable)

Exchange 2010 Persistence

Outlook MAPI / RPC Client Access Service

- Enable Static Ports for RPC CAS

Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSEExchangeRPC\ParametersSystem
Value: TCP/IP Port
Type: DWORD

- Specify the Static Ports on CAS (RTM)

- Microsoft.Exchange.AddressBook.Service.Exe.Config
- Set **RpcTcpPort** to the port value you want the Address Book Referral and NSPI interfaces to use
- Do NOT change **NspiHttpPort** or **RfrHttpPort**

- SP1: Set port in the registry

Key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSEExchangeAB\Parameters
Value: RpcTcpPort
Type: DWORD

- The setting is NOT migrated on upgrade – be sure to set it after upgrade

Exchange 2010 Persistence Outlook MAPI / RPC Client Access Service (SP1)

- Allow Kerberos for authentication to the CAS Array behind a Hardware Load Balancer
 - RTM was limited to NTLM – LSASS bottleneck
 - Need to create an AlternateServiceAccount for use
 - Script to assist with setting the service account with SP1
(RollAlternateServiceAccountCredential.ps1)

<http://technet.microsoft.com/en-us/library/ff808313.aspx>

Exchange 2010 Scaling

- Large scale Exchange deployments have unique NLB scaling issues
 - Devices support maximum throughput and/or connections
 - MAPI clients via NAT – no current solution for large scale (Source IP = same)
 - MAPI concentrators
 - Archiving solutions that “scrape/stub” via MAPI
 - BES versions that use MAPI
 - Single client concentrator -> single target CAS

Exchange 2010 Operations

- **Single VIP or Multiple VIPs?**
 - All protocols on a single VIP? iRule complexity
 - Multiple VIPs allows scale-out vs. scale-up
- **Consider using multiple VIPs when possible**
 - Additional stats / logging options for client traffic types
 - Segmenting client traffic
 - Reduce iRule complexity / better persistence

SSL Offloading

- SSL Offload vs. SSL Bridging
- When should SSL Offload be used?
 - Reduces SSL processing overhead on CAS
- When shouldn't it?
 - Security requires encryption from LB to CAS

SSL Offloading

How to Configure SSL Offloading

- Outlook Anywhere
 - Set-Outlook Anywhere –SSLOffloading:\$true
- OWA, ECP
 - Registry Key
 - HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\MSExchange OWA
 - New DWORD: SSLOffloaded
 - 1 to enable, 0 to disable
 - Uncheck “Require SSL” in IIS Manager

SSL Offloading

How to Configure SSL Offloading

- Exchange Web Services
 - Uncheck “Require SSL” in IIS Manager
 - Edit the web.config file in `<exchange install>\ClientAccess\exchweb\ews`
 - Change all occurrences of **httpsTransport** to **httpTransport**
 - Run `iisreset /noforce`
- Autodiscover
 - Uncheck “Require SSL” in IIS Manager
 - RTM: Edit **web.config** in `\ClientAccess\Autodiscover`
 - SP1: DON'T edit web.config!

Load Balancer Timeouts

- Rule of Thumb: One Hour
- OWA/ECP: As long as the longest session timeout (private or public profile)
- EWS: Based on the EWS application
- OA: Use the rule of thumb

Agenda

- History of Load Balancing
- Load Balancing Principles
- Exchange 2010 Load Balancing Requirements
- Hardware Load Balancing Vendors

Exchange 2010 SP1 Kemp LoadMaster

demo
Microsoft®
Exchange
CONNECTIONS

Load Balancing

Options and Pros/Cons

Requirements		Solution		
Fail-over	Capacity	Cost	Technology	Client<->CAS Affinity
Automatic. Minimal client downtime.	+++	\$\$\$	Hardware Load Balancer. E.g. F5 BigIP, Citrix NetScaler	Fall-back between Existing Cookie LB-created Cookie SSL ID Source IP depending on protocol/client
	++	\$\$	Software Load Balancer in separate server layer. E.g. TMG or UAG.	Selects either "LB-created Cookie" or "Source IP" depending on protocol/client
	+	\$	Software Load Balancer in same server layer as CAS. E.g. Windows Network Load Balancer (WNLB).	Source IP
Manual steps to detect issues and to fail over. Client DNS caches cause slow fail-over.	+++	\$	DNS Round Robin	Each client gets a random CAS IP address
	+	-	No Load Balancer	You manually assign separate hostnames for each CAS

Load Balancing Vendors

Vendor Name	Product Line	Website
A10	AX Series	http://www.a10networks.com/
Avanu / CAI Networks	Webmux	http://www.avanu.com/products/webmux.htm
Barracuda	Barracuda Load Balancer	http://www.barracudanetworks.com/ns/products/balancer_overview.php
Cisco	ACE (on switch/router)	http://www.cisco.com/en/US/products/ps6906/index.html
Cisco	CSS (dedicated)	http://www.cisco.com/en/US/products/hw/contnetw/ps792/
Citrix Systems	NetScaler	http://www.citrix.com/English/ps2/products/product.asp?contentID=21679
Brocade	ServerIron	http://www.brocade.com/sites/dotcom/products-solutions/products/ethernet-switches-routers/application-delivery/product-details/serveriron-adx-series/index.page
F5	Big IP Local Traffic Manager	http://www.f5.com/products/big-ip/
Kemp	LoadMaster	http://www.kemptechnologies.com/en/load-balancer.html
RadWare	AppDirector	http://www.radware.com/Products/ApplicationDelivery/AppDirector/default.aspx

Load Balancing Vendors

- Virtual Environments

- [Kemp Virtual LoadMaster](#) (VMWare)
- [Kemp Virtual LoadMaster](#) (Hyper-V)
- [F5 BIG-IP VirtualEdition](#) (VMWare and Hyper-V)
- [Citrix NetScaler VPX](#)

- Additional WIKI information

- <http://social.technet.microsoft.com/wiki/contents/articles/exchange-2010-client-access-array-amp-load-balancing-resources.aspx>
- Links to Exchange 2010 Guides
 - F5 BIG-IP, Barracuda, Citrix NetScaler, Brocade ServerIron, Kemp

Your Feedback is Important

Please fill out a session evaluation form
drop it off at the conference registration
desk.

Thank you!