

Securing & Assuring Networks for Converged Services

Interop NY - Sept 21, 2006
Gregory M. Lebovitz
gregory@juniper.net



- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



Measure the ROI of the Mechanism

- For Resiliency planning and security budgeting measure cost of loss of
 - the assets you are trying to protect, and/or
 - the services you are trying to deliver



- Where:

C_i is Cost to Implement

C_l is Cost if Loss Occurs

P is probability rate

t is a given time period (yr)

$$C_i(t) \leq C_l(t) * P(t)$$

- Conclusion: Implement if and only if:
Probable Cost of Loss > Implementation Cost for a fixed time period

- Measure your ROI for Resiliency
- **Speaker: Gregory Lebovitz**
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



Gregory Lebovitz

- Joined NetScreen March, 1998
- Office of CTO at time of acquisition in 2004
- Lead Security Solutions Engineering
- Technical Director & Architect
 - Advanced Technology Architect, Projects:
 - NSRP, Routing Protocols, Securing MCAST, IPv6, VoIP, RapidDeployment, IPsecVPNs, etc.
 - Strategist, M&A:
 - Neoteris (SSL-VPN), Kagoor (VoIP), Peribit (WAN Excel), Redline (App Front End, LB)
- Standards - IETF
 - PKI4IPsec WG Co-Chair


- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- **Interoperability**
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



Converged Networks: Interoperability is your Biggest Challenge



- Compatibility of the elements at the app layer
- Current state of affairs
 - Read Interop Reports, SIPit, Go to bake-offs



- Try before you buy, get in the lab & test before you deploy

- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- **Intra-Device Resiliency**
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



Consider 4 Areas of Intra-Device Resiliency


- OS Architecture
- Processor Redundancy
- Software Features to Ease Human Error
- Link Layer Redundancy

Consider 4 Areas of Intra-Device Resiliency

- OS Architecture
- Software Features to Ease Human Error
- Processor Redundancy
- Link Layer Redundancy


Networking OS Engineered for Resiliency

Dedicated Resources



- Reduces service outages when unforeseen events occur
- Predictable performance for voice, video and other time critical apps
- Modularity for full router control while under attack
- Add 1000's of filter terms w/o perf. degradation

Modular Software



- Minor problems do not lead to system crashes
- Next Gen CLI prevents operator error
- Graceful restart and M10i hitless switchover
- VRRP, APS, MPLS FRR

Multiple Modes of Multitasking


Cooperative

&

Pre-Emptive

- Multiple scheduler types are necessary
- Cooperative is Run-to-Completion; e.g. Routing DB recalculation and route entry
- Real Time is Pre-Emptive, "I need the processor NOW!"; e.g. hellos, adjacency message

Single Release Train

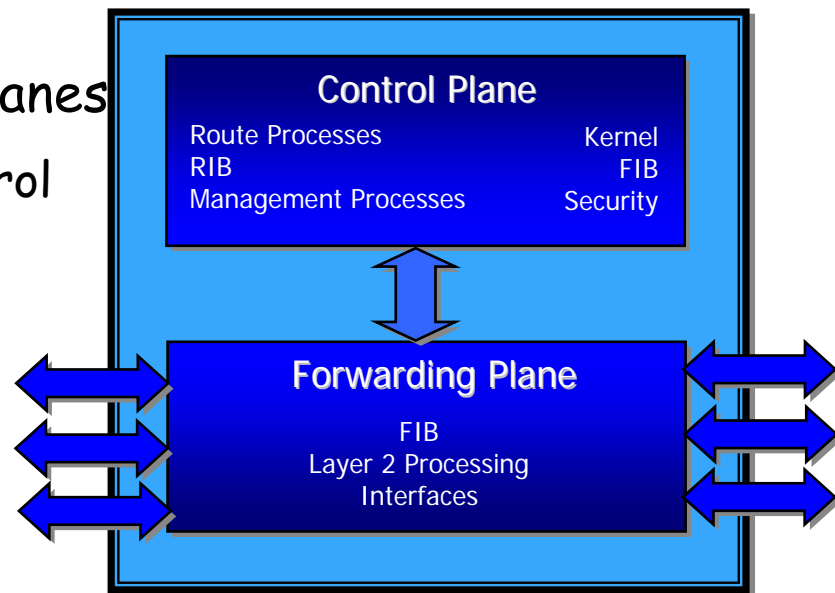



- Structured quarterly release process
- One release across all platforms with all features included
- Reduced operational burden, provides real choice

Juniper your Net

Prerequisite Elements for High Control Plane Availability

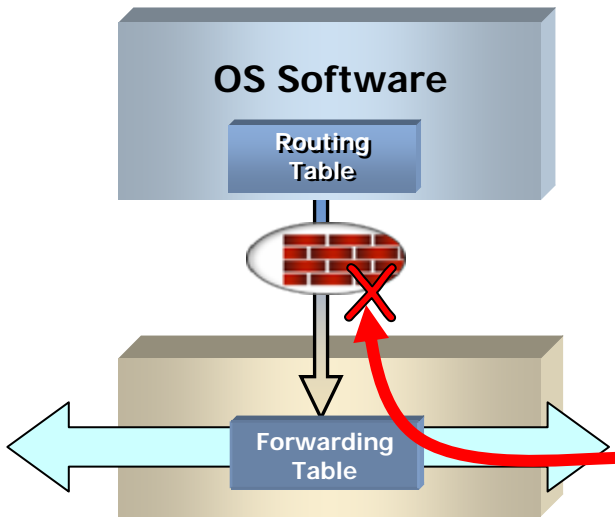

- Separation of control and forwarding planes
 - Allows continuing forwarding during control plane restarts



- Stateful control plane redundancy
 - Enables control plane recovery from unplanned outages
 - Enables graceful switchover for planned and unplanned outages
 - Maintenance
 - Software upgrades

Dedicated Resources

Industry's Most Secure Routers



- System resources reserved ensuring full router control while under attack (DoS)
- Stateful firewall, rate limit protection for control plane
- Critical filters can be added/changed dynamically while under attack; Thousands of filters at line rate
- Forwarding engine free to forward traffic when routing engine is swamped (say topology is flapping back and forth)

Juniper your Net

Consider 4 Areas of Intra-Device Resiliency

- OS Architecture
- Processor Redundancy
- Software Features to Ease Human Error
- Link Layer Redundancy

The Problem of Human Error

- Most network downtime is due to NOT to software or hardware failure, but to **HUMAN ERROR**
- 99.999% uptime impossible unless human error sharply reduced

"The possibility of failures would be much reduced if you consider that changing device configuration causes 60% of downtime due to human error."

-Jeffrey Nudler, Senior Analyst

<http://networkworld.com/news/2005/101005-ietf.html>

Features Required for Minimizing Operator Error

- Hierarchical configuration files
 - Simplifies interpretation and management
- Candidate configurations
 - Enables careful inspection before committing
- Explicit commit
 - Changes become active only when the operator is ready, or at a predefined commit time
- Rollback Configs stored locally
 - Quick and easy backouts when things go wrong
- Extensive error checking & Commit Scripts

Example: Error Checking

- Check if IGP is configured for all `so-*` interfaces
- Check if Ingress RPF filters enabled for all `so-*` interfaces

...and optionally, correct the error!

Example: Enforce Configuration Rules

- "Maximum number of VLANs per port"
- "All public exchange peers must have MD5"
- "Firewall has to have trailing explicit deny with logging enabled" (helps catch the goobers)
- "Certain set of parameters must always be set"
 - Filter on lo0, telnet disabled
 - Export filters on ebgp sessions

...commit prevented until configuration
is in compliance with rules.

Script Macros


- Scripts are written and tested by tier-3 engineers
- Scripts are uploaded to routers
- Operators invoke macro and enter variables
- Macro writes the configuration
 - Complex configurations can be created from simple entries
- Result:
 - Simple operator entries sharply reduce configuration errors
 - Configurations written correctly **every time**
 - Configurations written consistently **every time**
- Available in JunOS since 7.4

Consider 4 Areas of Intra-Device Resiliency

- OS Architecture
- Processor Redundancy
- Software Features to Ease Human Error
- Link Layer Redundancy

Multiple Physical Interfaces

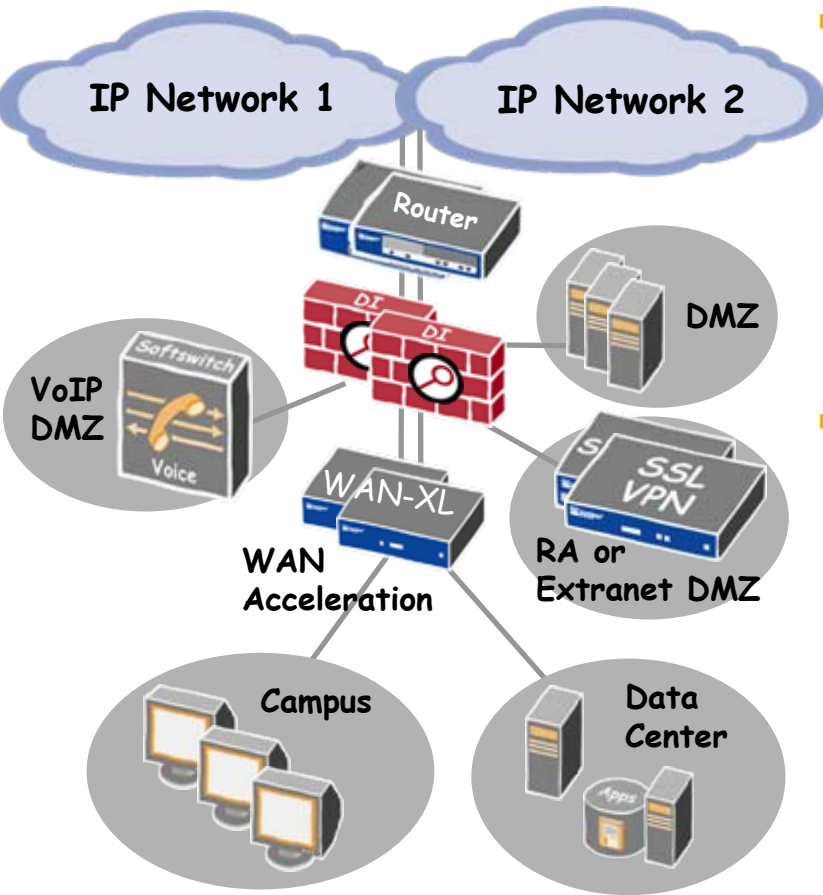
- L2
 - WAN Links
 - One higher bandwidth one lower bandwidth
 - Ethernet
 - STP, or link-up-non-forwarding
- L3 - 2 router interfaces to upstream peers
- 2 HA links (if device takes dedicated HA ports)



- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- **Resiliency in Network Design**
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



The Importance of Network Segmentation



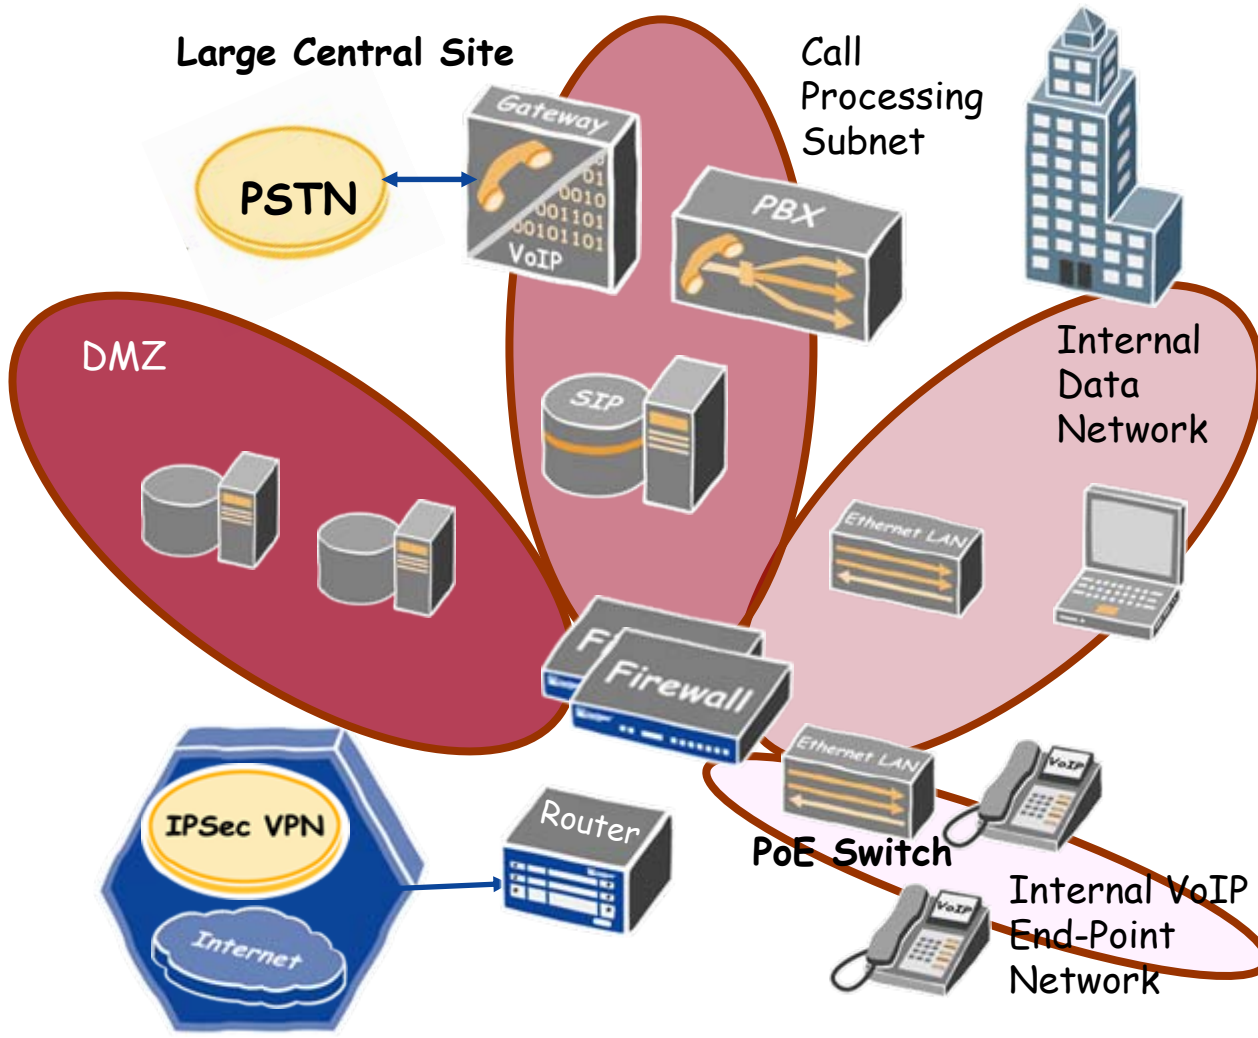
■ WAN Gateway Challenges

- Maximize availability, resiliency, quality
- Protect public facing servers and infrastructure
- Optimal support for broad mix of app & traffic
- Massive # VPN Connections or Large BW single tunnels

■ WAN Gateway Solutions

- High performance Enterprise routers dual connected
- MPLS for improved quality and traffic engineering
- High performance firewall/VPN, security gateway
- Intrusion Prevention mitigates threats
- SSL/IPsec VPN Gateway for secure access for remote workers and partners
- WAN Optimization to remote locations

Network Isolation - VoIP Example



Redundant security devices for failover and high availability

Router - QoS policy and scheduling

Scalable VPN supporting thousands of connections

ALG technology to extend corporate VoIP

MPLS TE passed to provider MPLS network

Zone architecture for intra/inter zones with policy enforcement

- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- **Stateless HA**
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



Routing Protocols with Advanced Mechanisms

- OSPF, BGP, RIP, etc.
- Delicate balance in setting hello/adjacency timers
 - Too High vs Too Low
- Bi-Directional Forwarding Detection (BFD)
 - Below routing protocols
 - Essentially a neighbor monitor for any applicability
- Graceful Restart
 - Restarting Router and Helper Router
 - Grace period timers

In Service Software Upgrades - ISSU

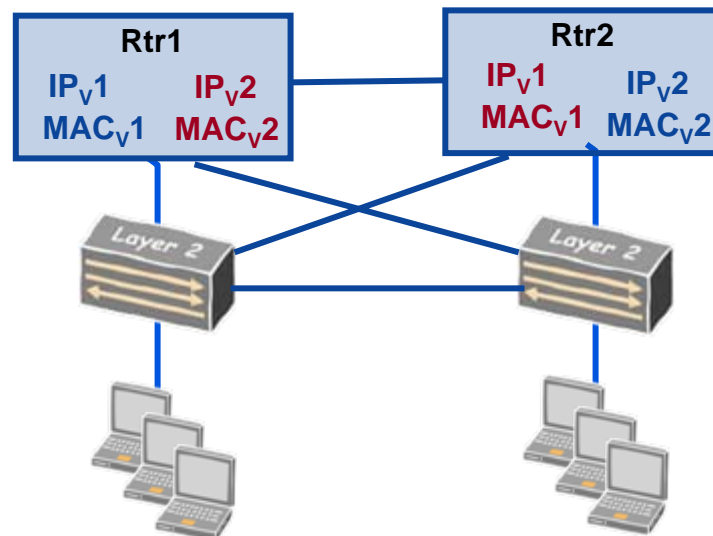
ISSU:

The seamless migration from
one release to another

- The *right* way to add new features
- ISSU enablers:
 - Control and forwarding plane separation
 - Graceful Routing Engine Switchover (GRES)
 - Non-Stop Routing (NSR)
 - And, yes, modular architectures

Forwarding Interface Availability

- VRRP & **RP protocols
 - Provide gateway failover for hosts
 - Standard and proprietary versions
 - VRRP
 - NSRP - Juniper NetScreen
 - **RP from some other vendors
- VRRP provides standardized gateway failover
 - Uses virtual IP to present single gateway to up/downstream hosts
 - Uses virtual MAC to signal FDB changes on switches for seamless failover from end host perspective
 - Provides quick and easy failover
 - Can operate in A/P as well as A/A environments via multiple "groups"
 - Fail-over is "one sided"
- VRRP works well where no state sharing mechanisms needed: pure packet forwarding, hosts to default gateway



- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- **Stateful HA**
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



NSRP for Stateful Failover


Feature	NSRP	VRRP
Virtual IP and MAC Addresses	✓	✓
Multiple Virtual Interfaces per device	✓	✓
Protocol Layer	Layer 2, Ethernet	Layer 3, IP
Terminate Packets by non-Interface “owner”	✓	
Master Pre-election	✓	
Sub-second fail-over	✓	
Preserve Connection during fail-over (stateful fail-over)	✓	
Path Monitoring (Link check, TrackIP)	✓	
Mirroring of Run-Time Objects	✓	
Secured protocol (encryption, authentication...)	✓	Extensible, not standard

- VRRP RFC 3768
- Path Monitoring mechanisms for detecting outages are critical:
 - Link, neighbor processor, two hops beyond

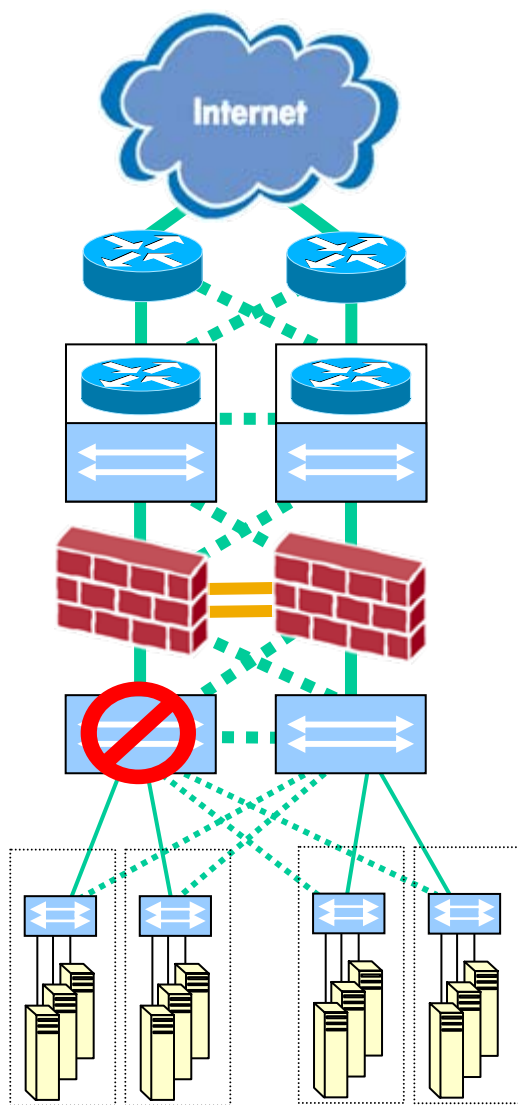
Redundant Matrix Step-by-Step

■ Upstream Switch Fails

- Detect Failure - Activate Secondary "Untrust" Link
- Traffic forwarded via secondary "Untrust" interface on affected Firewall/VPN device
- Throughput maintained - assuming sufficient switch capacity - Firewall/VPN Fail-over unnecessary




Redundant Matrix Step-by-Step



■ Downstream Switch Fails

- Detect Failure - Activate Secondary "Trust" Link
- Traffic forwarded via secondary "Trust" interface on affected Firewall/VPN device
- Throughput maintained - assuming sufficient switch capacity - Firewall/VPN Fail-over unnecessary

Redundant Matrix Step-by-Step



- Firewall/VPN device fails

- Fail-over happens, all active sessions, VPN tunnels, Security Associations maintained
- Sub Second Fail-over
- Throughput reduced

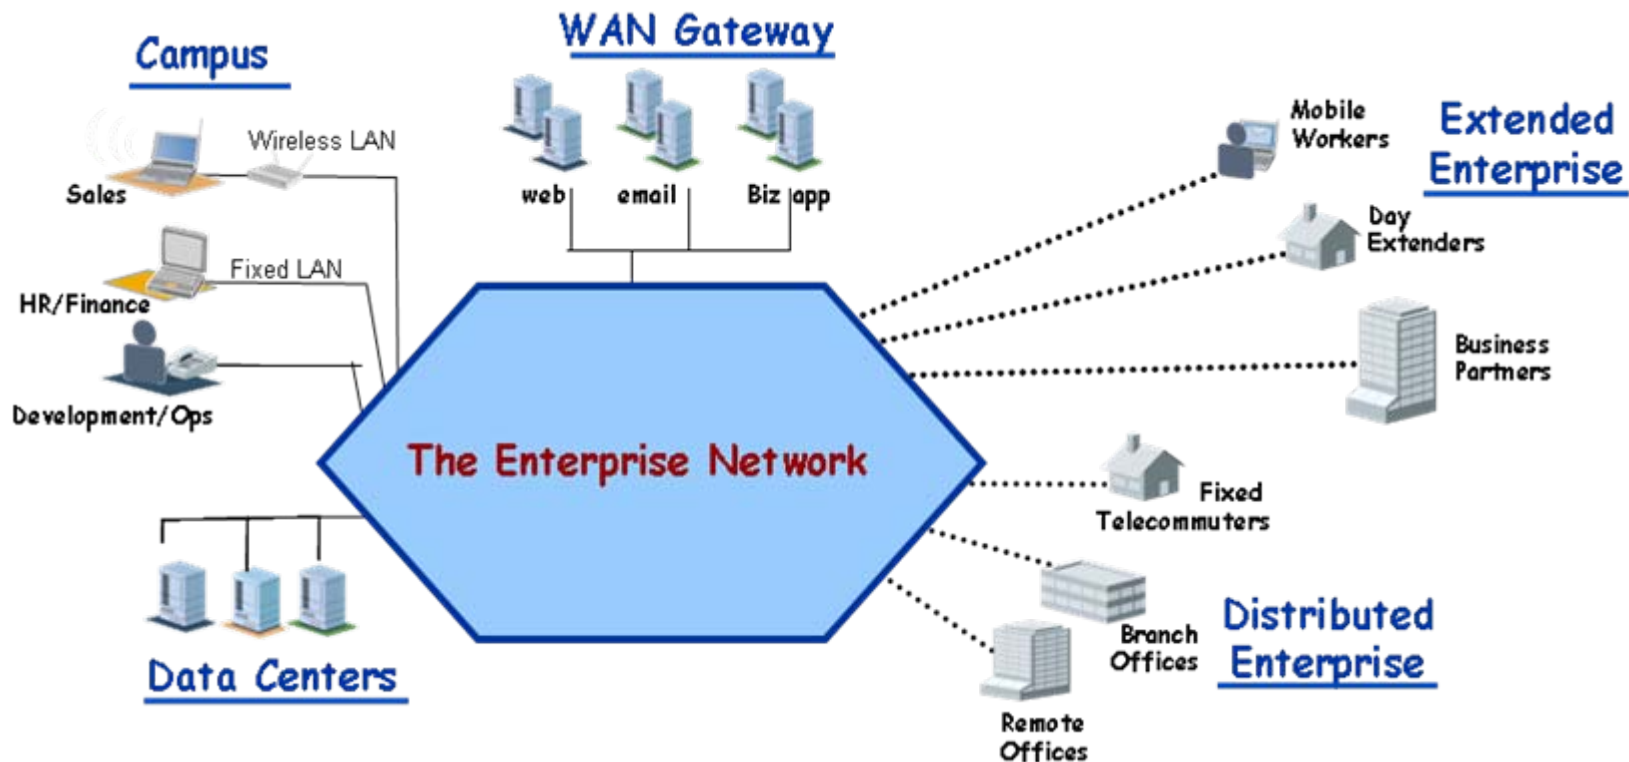
NSRP Conclusion

- Provides intrinsic support for all firewall and IPSEC failover in environment around the box.
- State sync and config sync are part of the protocol, so no need for anything extra.
- Active/Passive and Active/Active modes supported.
 - A/P tends to be a simpler design as A/A requires load balancers or different hosts to point to different groups for load sharing.

- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- **Branches Need UTM**
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



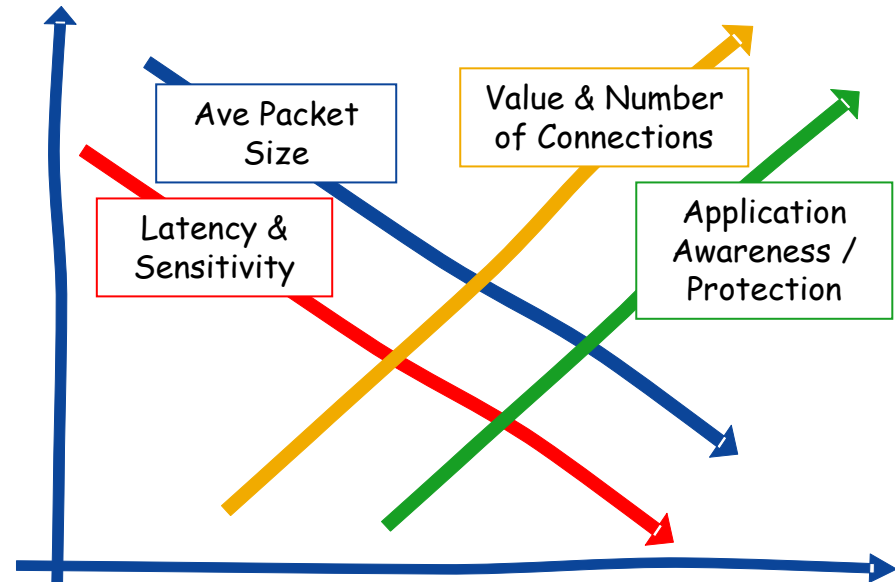
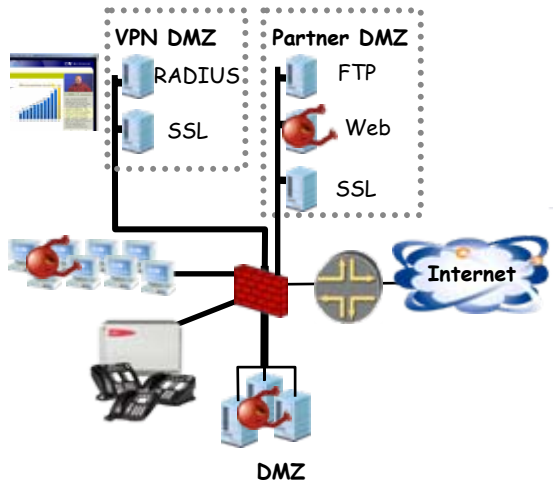
Evolving Challenges and Requirements



Need a secure and resilient infrastructure able to deliver differentiated applications and services across the network

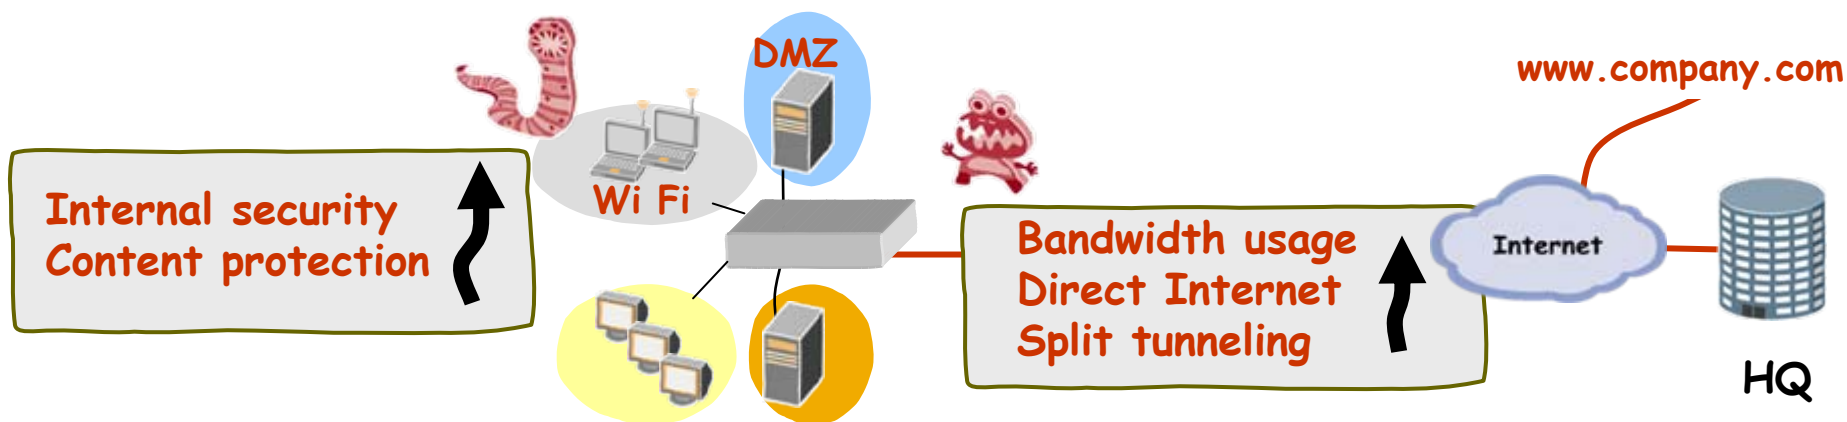
- Single IP infrastructure - demanding applications require network performance
- Virtual enterprises - dynamic perimeters, different users, devices, locations and trust levels
- Elevated threat environment - application level attacks and worm propagation
 - Regulatory compliance (now global) - granular access controls and auditing

WAN Gateway Requirements



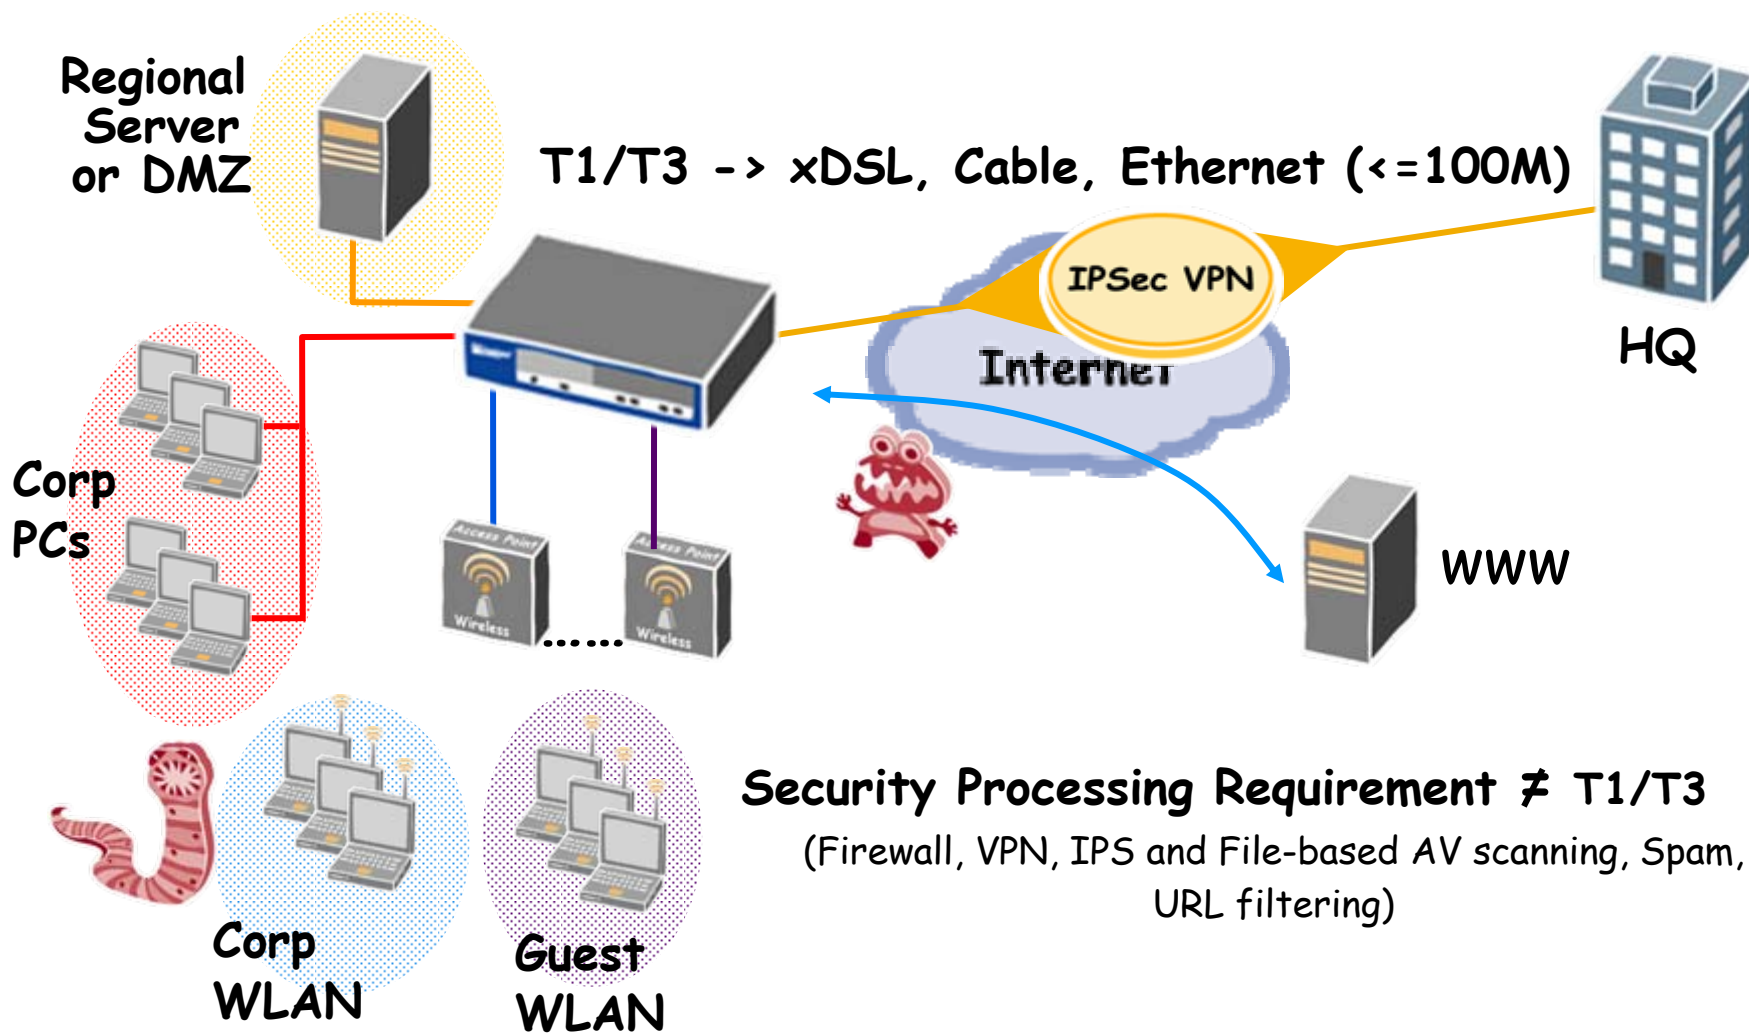
- Provide high performance for large and small packet traffic mix
- **Make traffic decisions with low latency to ensure applications are not affected**
- **Handle traffic load, complexity & availability requirements as # & value of connections increase**
- **Understand application requirements and prevent/mitigate application-level attacks**

Regional/Branch Office Trends



- Increased migration towards the branch/remote offices (from ~85% of employees in the branch in 2003 to 91% today) - Nemertes Research
- By 2007, 50% of the companies surveyed will significantly increase their WAN access bandwidth - Infonetics
- In 2005, 56% of companies had at least 1 internal attack
 - 65% had at least 1 external attack - CSI/FBI 2005 survey

UTM - Demands for Security Processing in the Branch



Security Processing Requirement \neq T1/T3
(Firewall, VPN, IPS and File-based AV scanning, Spam, URL filtering)

- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- **Bandwidth & Latency Tuning**
- Protect the Infrastructure from the Users



Over Provision, Use QoS for Emergencies

- Bandwidth is cheap, and getting cheaper
- Cut a Deal - ISPs are hungry for revenue
- Over provision for tomorrow
- Use QoS for Emergencies

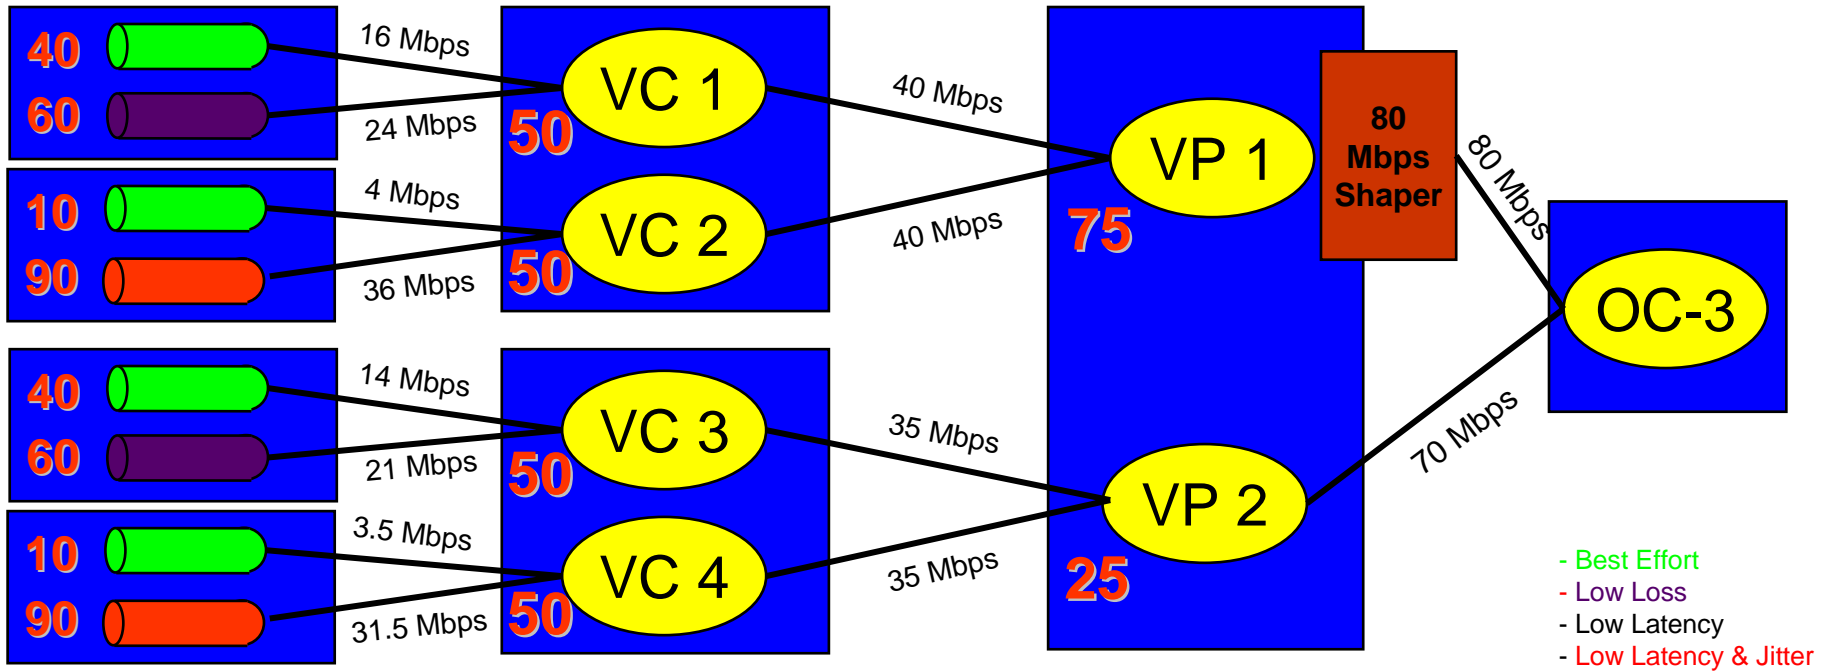


My Best QoS Advise: KISS



- Ingress Limiting
- Policing
- Shaping / Queueing - reservation, prioritization
- Number of Queues you need at any one point?
 - 4 (8 if you like to walk on the wild side)
 - Voice, Video, PriorityData, BestEffortData
- Hierarchical Queueing is a reality

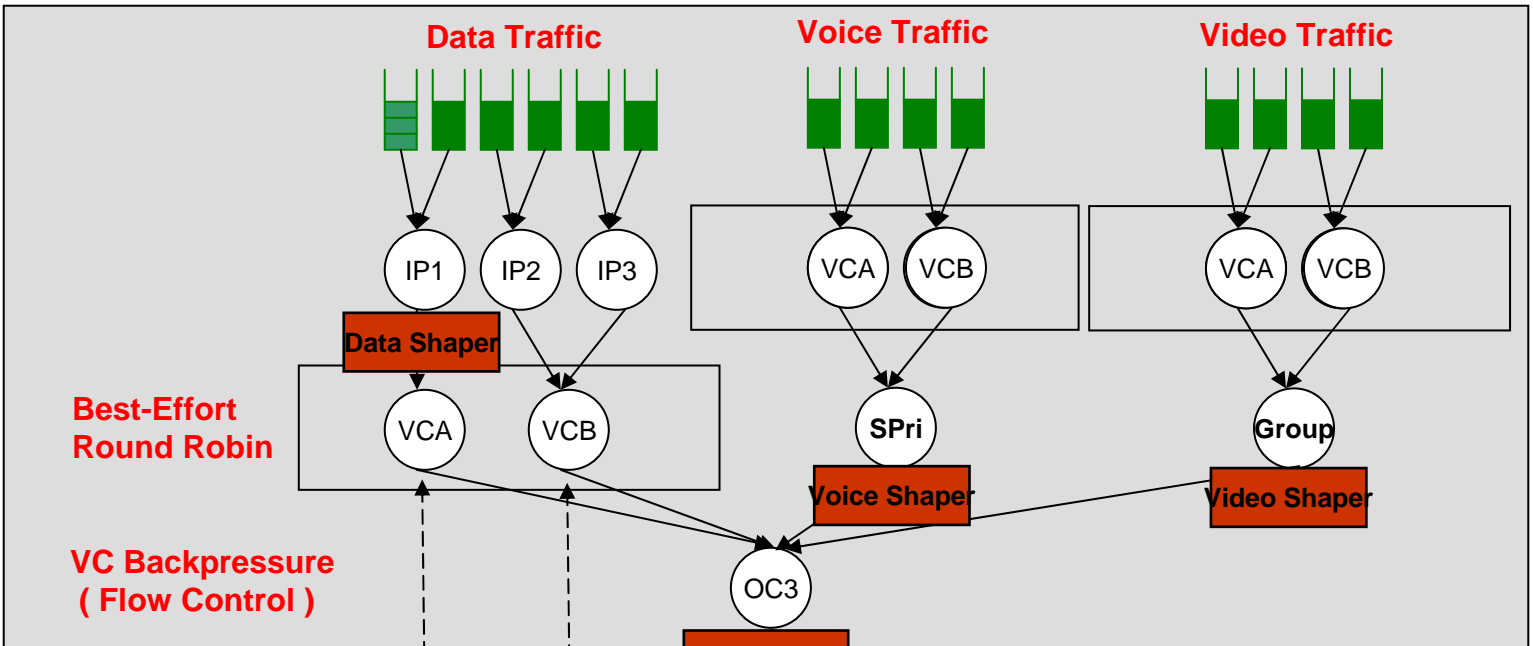
Hierarchical Rate Shaping



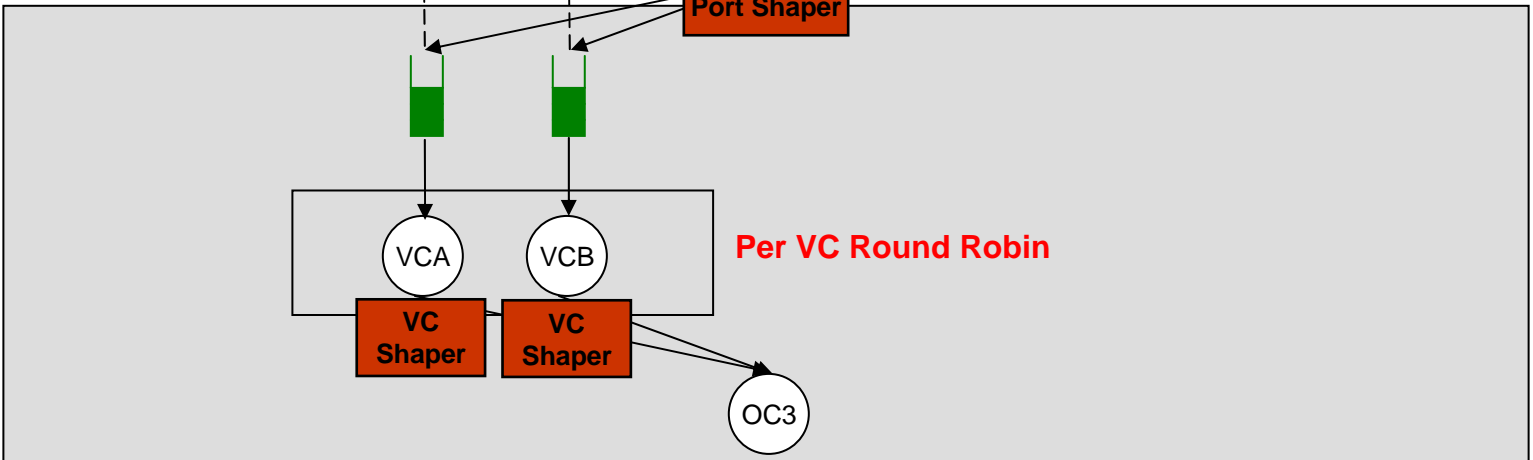
- Throttles the Node or Queue to a Max Bandwidth
- Excess packets buffered, not dropped. Reduces jitter.
- Can shape each interface and each queue
- Can reserve BW for other interfaces (starvation avoidance)

QoS in Converged Networks - Deployment Example

HRR Scheduler



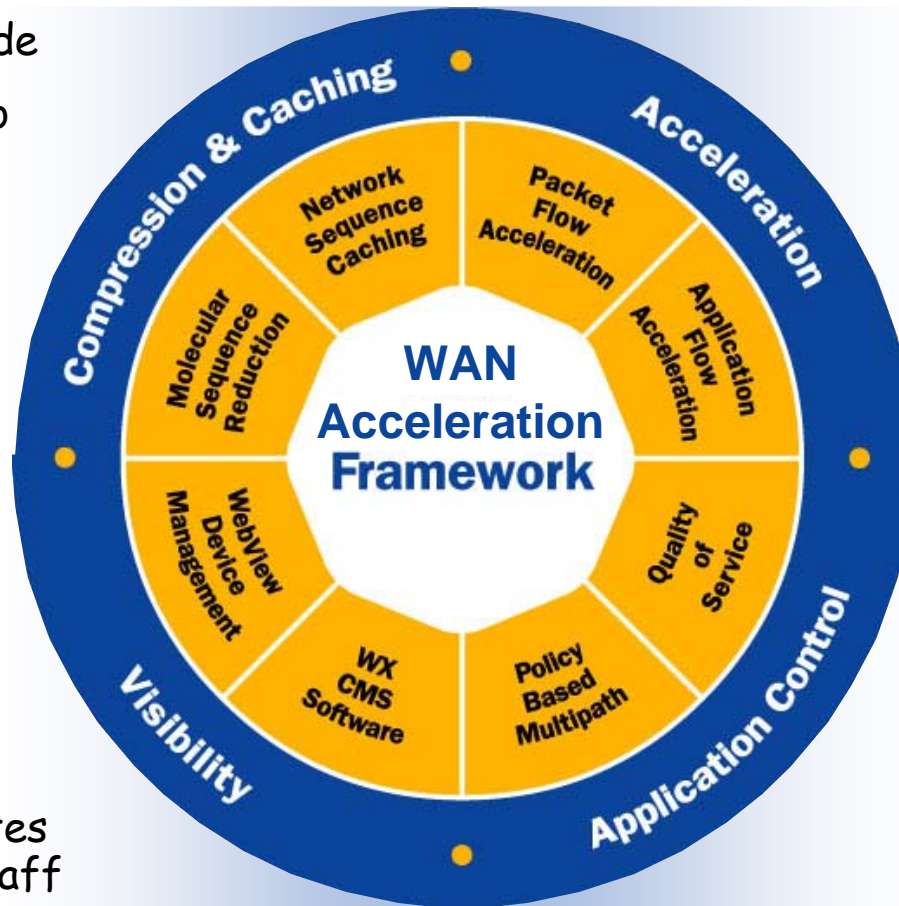
SAR Scheduler



WAN Acceleration Techniques

- Avoid WAN upgrade
- Increase access to shared resources
- Improve disaster recovery
- Enable new application rollout


- Reduce time to troubleshoot
- Automate remote device deployment
- Manage remote sites with no local IT staff



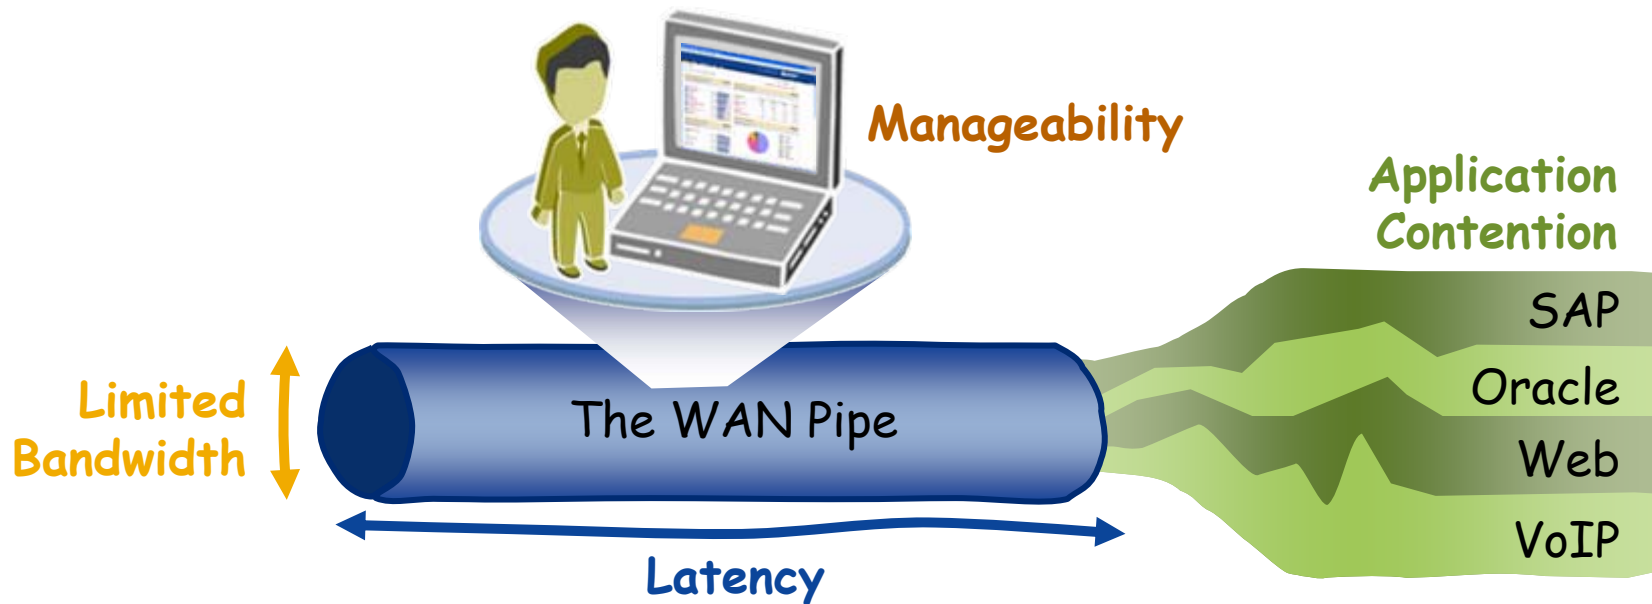
- Improve user productivity
- Centralize servers
- Consolidate data centers
- Achieve regulatory compliance

- Enable VoIP, other money-saving apps
- Ensure critical apps get priority service
- Make use of cheaper secondary links

WAN Acceleration Reference Architecture



Accelerating Applications over the WAN



Compression, Caching

Acceleration

Application Control


Visibility and Reporting

Juniper your Net

- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- **Protect the Infrastructure from the Users**

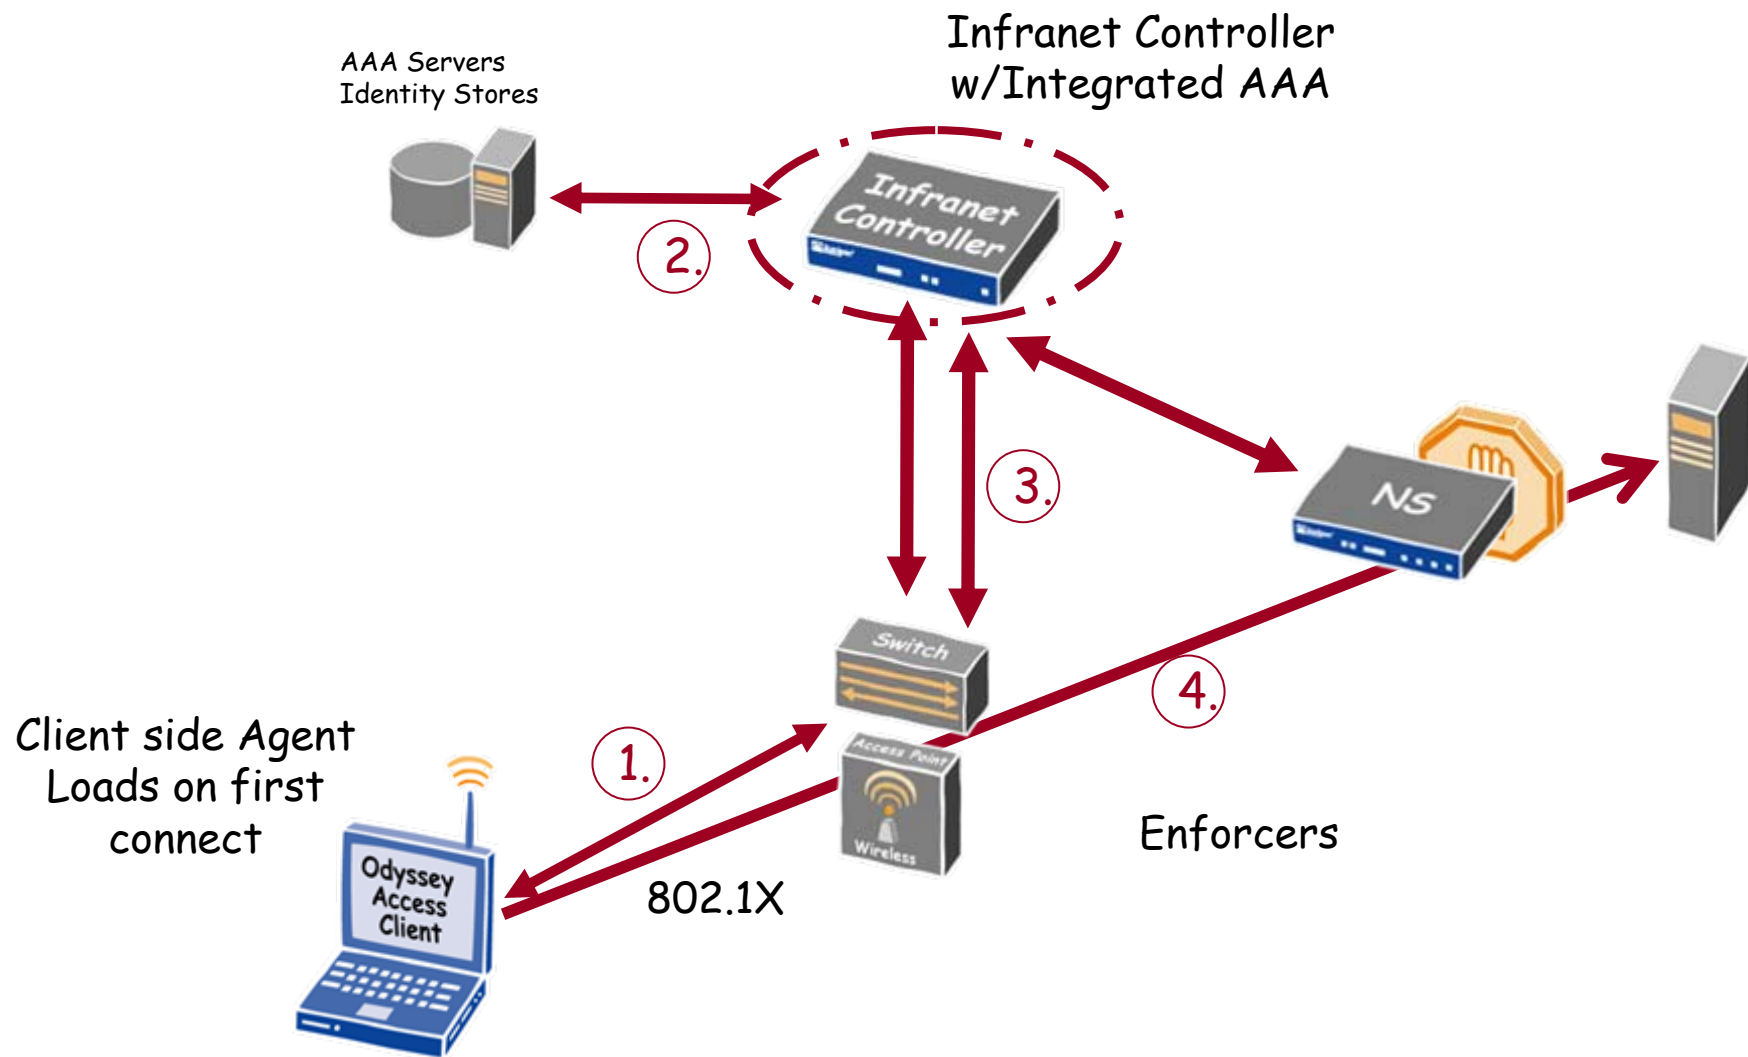


Enterprise Infranet - Unified Access Control Solution



UAC 2.0: Layer 2 + Layer 3

The future of Unified Access Control



- Measure your ROI for Resiliency
- Speaker: Gregory Lebovitz
- Interoperability
- Intra-Device Resiliency
- Resiliency in Network Design
- Stateless HA
- Stateful HA
- Branches Need UTM
- Bandwidth & Latency Tuning
- Protect the Infrastructure from the Users



Gregory M. Lebovitz
gregory@juniper.net

Juniper
your
Net™

