Securely Architecting Your Application for the Cloud

Alex Stamos

CTO

iSEC PARTNERS

Takeaways

- Current conventional wisdom on cloud computing is missing the point
- You cannot securely move into the cloud without re-writing your software
- Secure cloud applications “collapse the perimeter”
- Properly going through this process should leave you more secure than before

Convention Wisdom

- Best reflected in work of:
  - Cloud Security Alliance
  - ENISA Cloud Computing Report
  - NIST SWG Cloud Working Group
- Lots of focus on multi-tenant risk, little focus on operational changes
- Still worth reading, but somewhat reflect application of 2003 compliance standards to new paradigm

Keeping It Real

- What are the realistic threats to cloud computing?

1. Loss of credentials via attacks against individuals
    - Spear-Phishing, malware, rubber hose
    - Gain access to (under EC2):
        - List of machines
        - Persistent Storage (EBS, SDB and S3)
        - Consoles
    - Don't automatically get access to:
        - Running machine state/memory
        - Login credentials
        - Non-persistent storage

Keeping It Real

2. Operational security breakdown
   - Going from 50 machines/sysadmin to 500 is life-changing
   - Need to plan from the start your security process
     - Patching
     - Hardening
     - Identity management
     - Logging
     - Application identification
     - Distribution of secure files
     - Forensics and IR
   - This is where a direct port to the cloud kills you

Keeping It Real

3. Misuse of new cloud technologies
   - Security promises of new technologies aren’t well understood
     - i.e. Access control in Hadoop
   - Easy to poorly architect system
   - Easy to downgrade security via change
     - Security zones in AWS
     - vShield zones in VMW based cloud

Silly Concerns

- Physical datacenter security
- Side-channel attacks

Not So Silly Concerns

- Xen 0-days
- Attacks against Intel VT-x instructions
- “Cloud cartography”
- APT Against Amazon Corporate

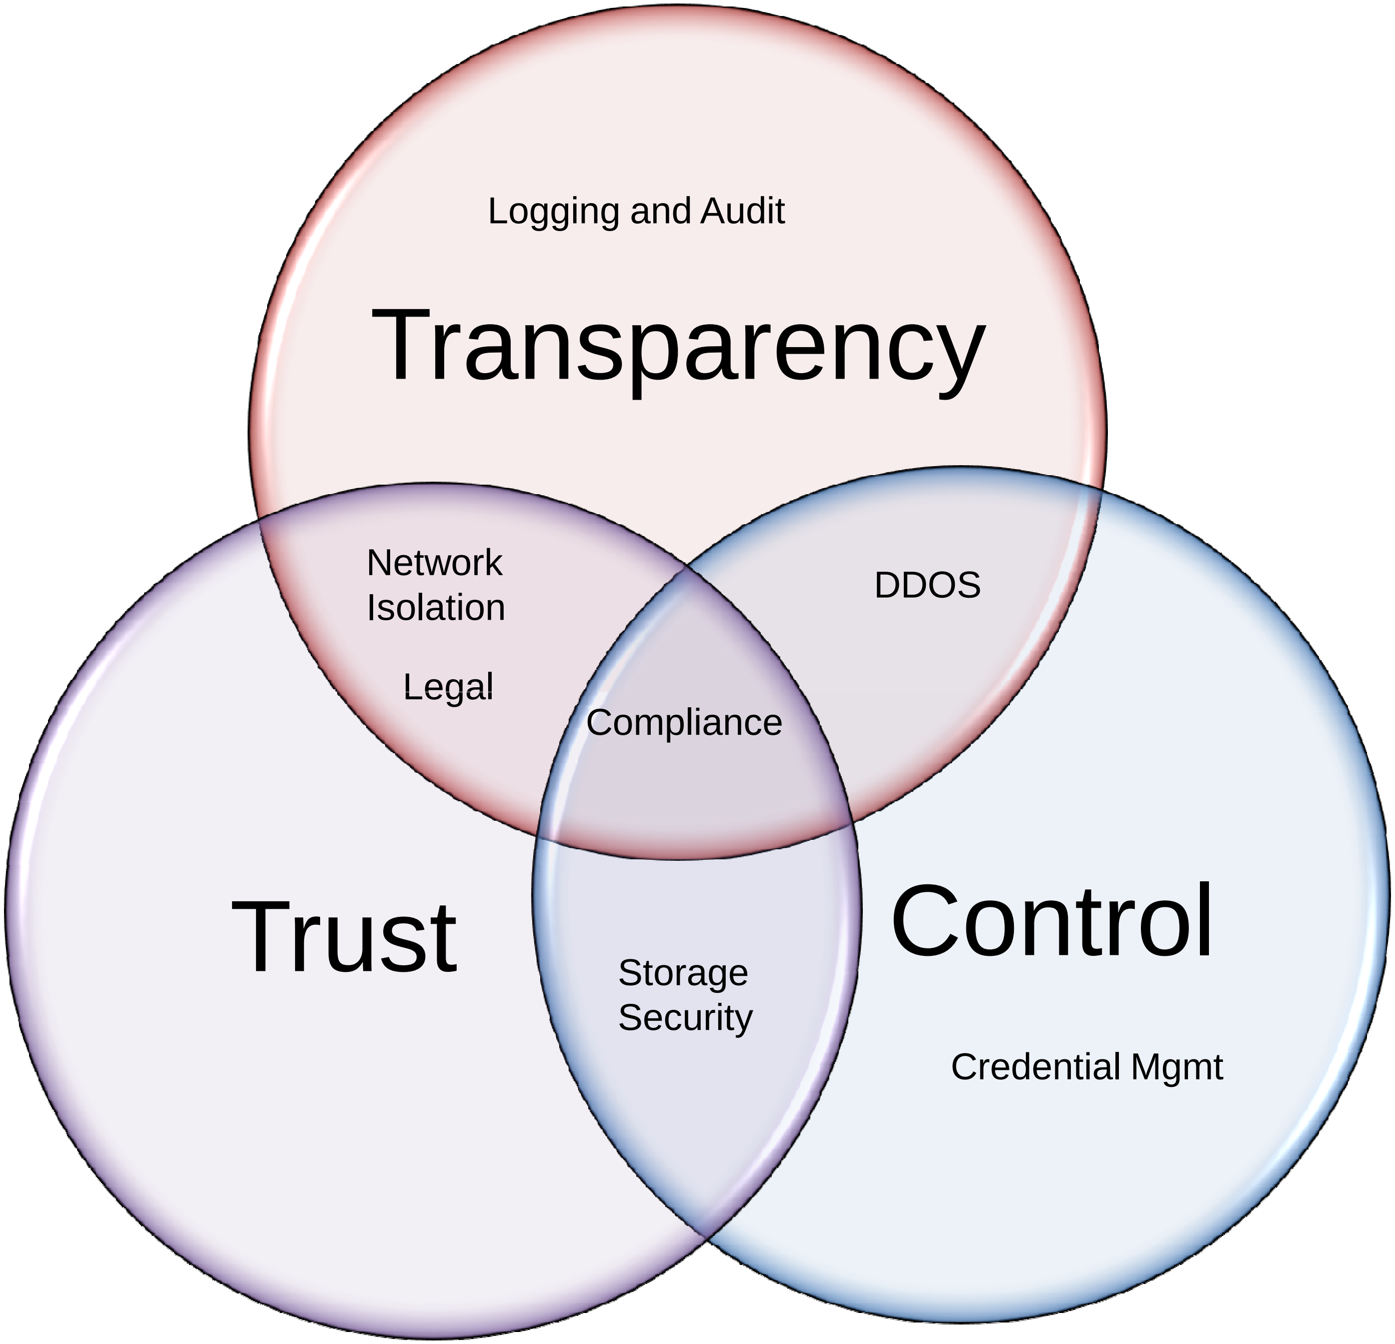
Logging and Audit
Transparency
Network Isolation
Legal
DDOS
Compliance
Trust
Control
Storage Security
Credential Mgmt

Why build on the cloud?

Efficiency is Unbeatable

- Cloud and virtualization technologies are unstoppable from a power/cost efficiency standpoint

PUE Components: Typical vs. Google

Impact of power & cooling innovations and best practices

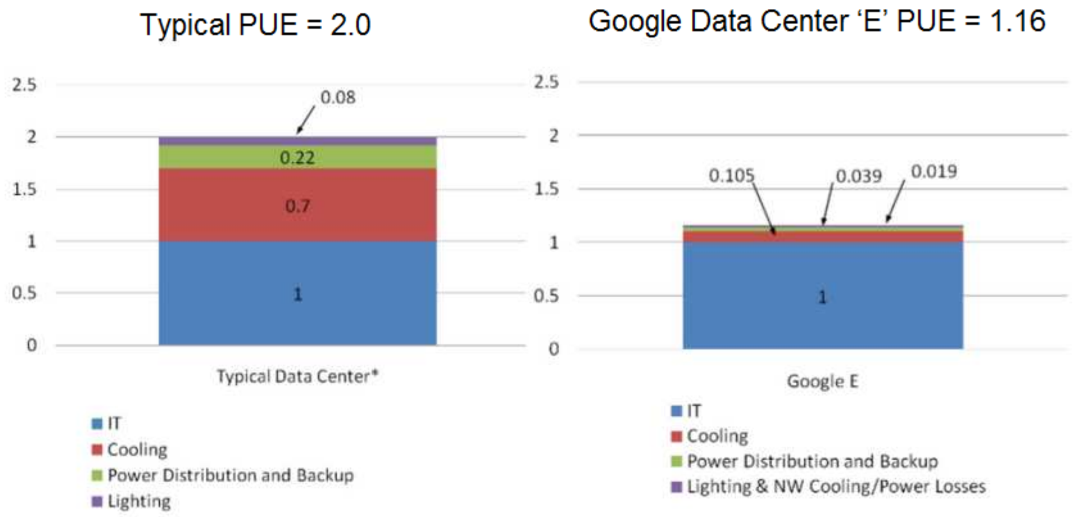

*Reference: Silicon Valley Leadership Group, Data Center Energy Forecast, Final Report July, 2008
Google E Data Center energy-weighted average PUE results from Q2-Q1'09 (to 3/15/09)

http://www.google.com/corporate/green/datacenters/summit.html

Custom Hardware...

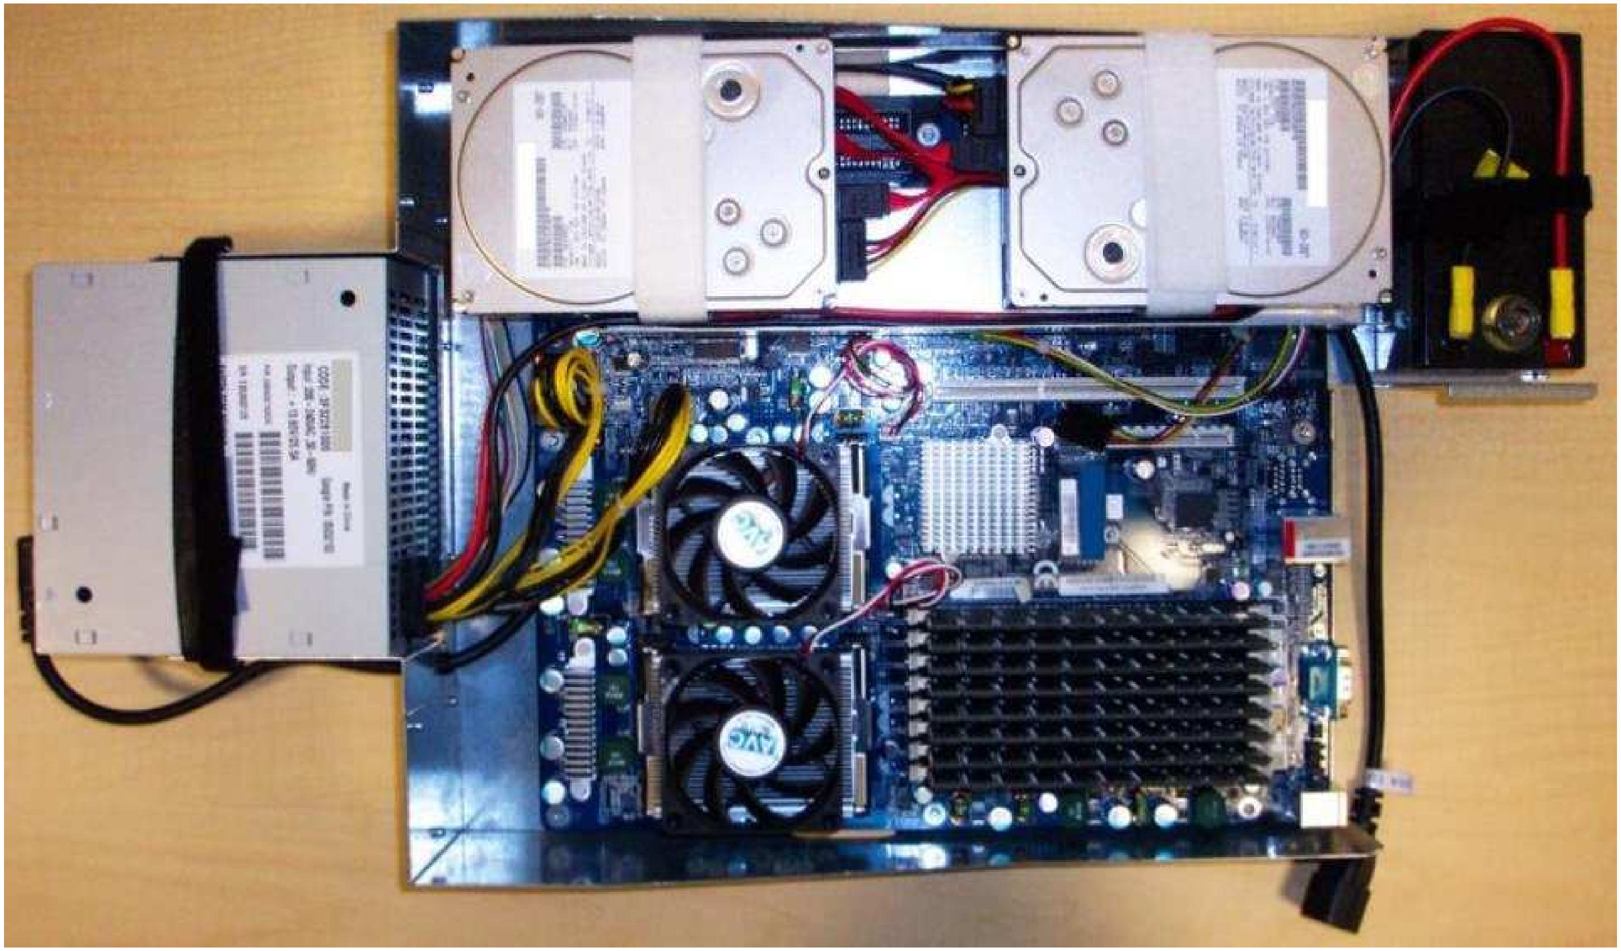

iSEC PARTNERS

…is becoming commoditized

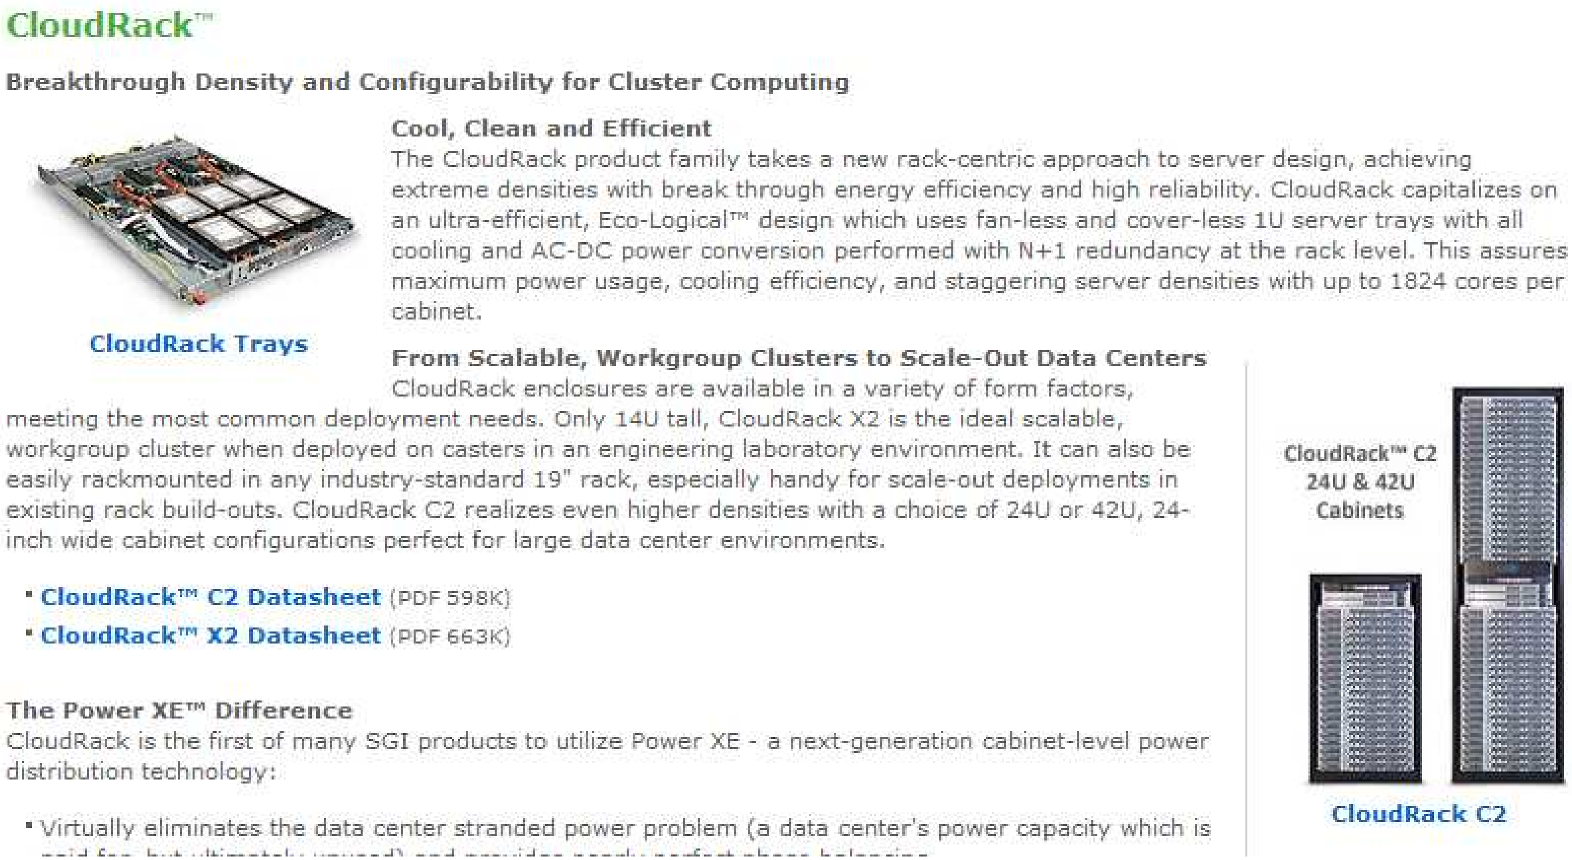

CloudRack™

Breakthrough Density and Configurability for Cluster Computing

CloudRack Trays

Cool, Clean and Efficient

The CloudRack product family takes a new rack-centric approach to server design, achieving extreme densities with break through energy efficiency and high reliability. CloudRack capitalizes on an ultra-efficient, Eco-Logical™ design which uses fan-less and cover-less 1U server trays with all cooling and AC-DC power conversion performed with N+1 redundancy at the rack level. This assures maximum power usage, cooling efficiency, and staggering server densities with up to 1824 cores per cabinet.

From Scalable, Workgroup Clusters to Scale-Out Data Centers

CloudRack enclosures are available in a variety of form factors, meeting the most common deployment needs. Only 14U tall, CloudRack X2 is the ideal scalable, workgroup cluster when deployed on casters in an engineering laboratory environment. It can also be easily rackmounted in any industry-standard 19" rack, especially handy for scale-out deployments in existing rack build-outs. CloudRack C2 realizes even higher densities with a choice of 24U or 42U, 24-inch wide cabinet configurations perfect for large data center environments.

- **CloudRack™ C2 Datasheet** (PDF 598K)
- **CloudRack™ X2 Datasheet** (PDF 663K)

CloudRack™ C2
24U & 42U
Cabinets

CloudRack C2

The Power XE™ Difference

CloudRack is the first of many SGI products to utilize Power XE - a next-generation cabinet-level power distribution technology:

- Virtually eliminates the data center stranded power problem (a data center's power capacity which is paid for, but ultimately unused) and provides nearly perfect phase balancing

iSEC PARTNERS

You too can own a box o' servers!

The future datacenter will have less:

- Centralized fault-tolerant storage
  - Clarions are the least cost effective devices ever purchased by IT
- Dedicated, specialized servers
  - 64 core DB server
- Dedicated, specialized network equipment
  - Lots of cheap boxes in my critical path
  - Create physical networking complexity, reduce flexibility
- Generic abstraction layer between DC and hardware
  - Standard square hole rack, 110v, cold/hot aisles

Before You Move, Ask…

- Why are you moving into the cloud?
  - Performance
  - Cost
  - Security

- What security promises are you making?

"Users will not be able to access the medical records of patients belonging to other medical institutions."

  - Access controls in web logic
  - Encryption of PII with per-doctor key
  - Secure coding

Traditional 3-Tier Architecture

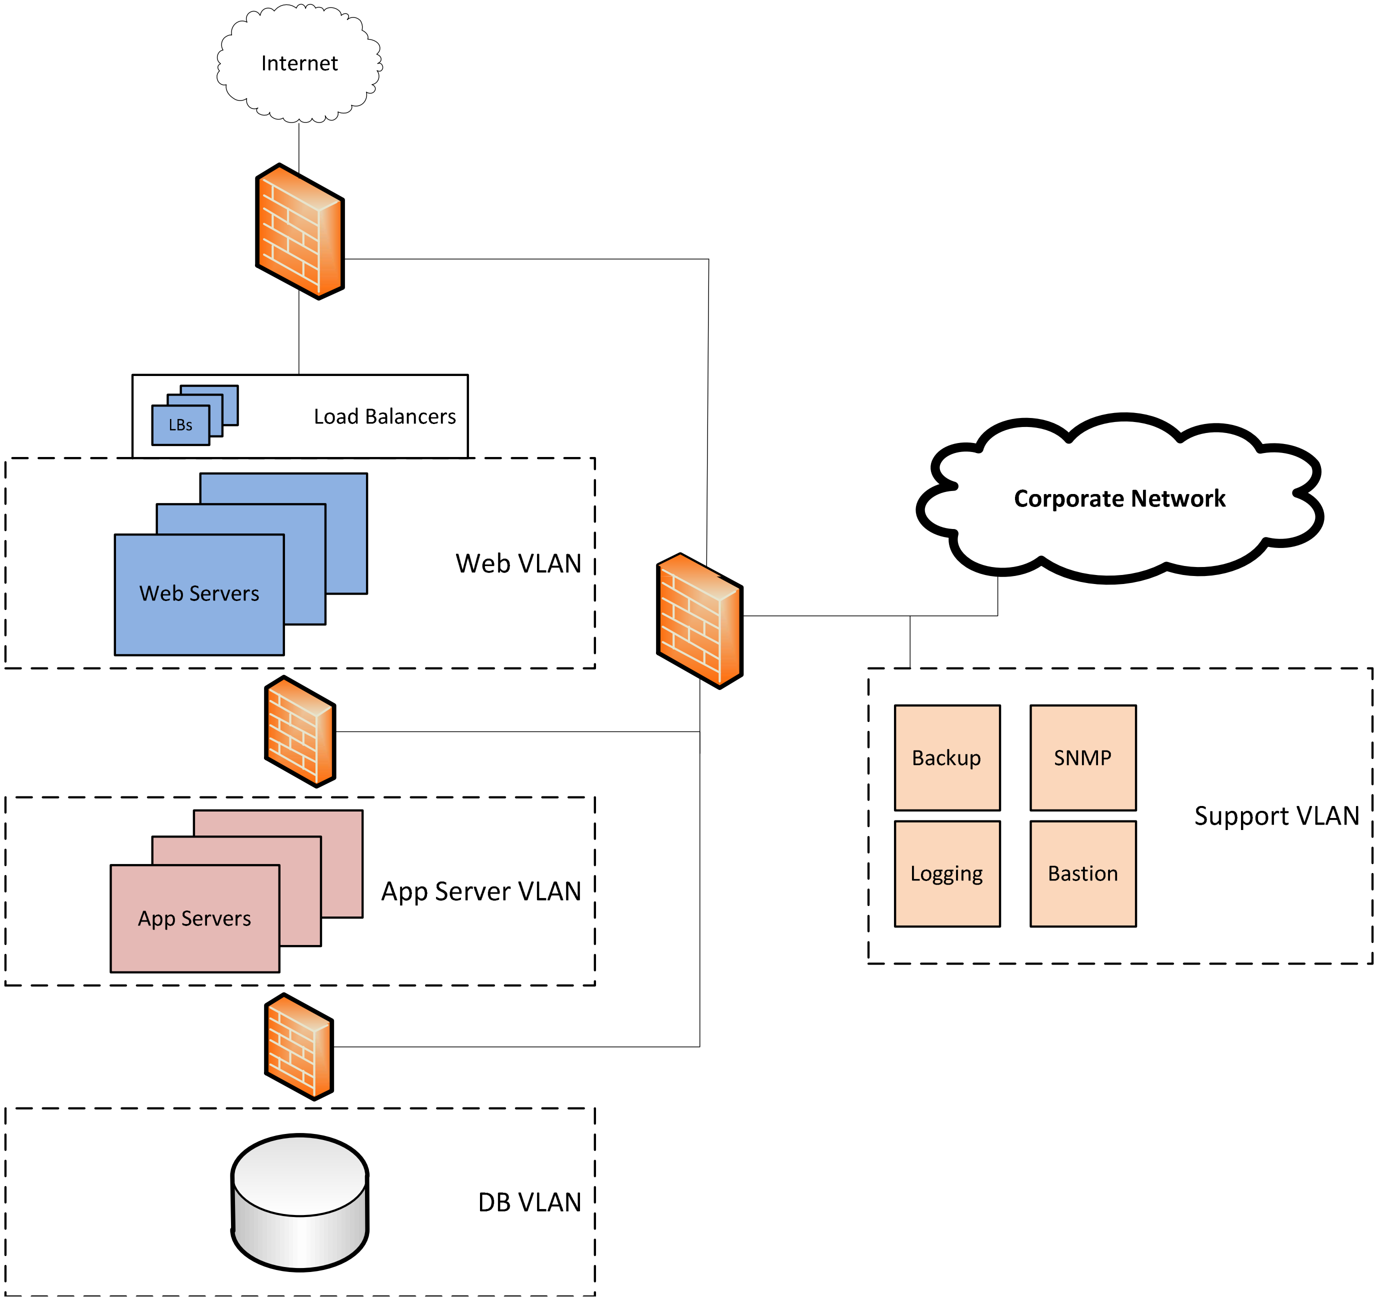

Physical configuration

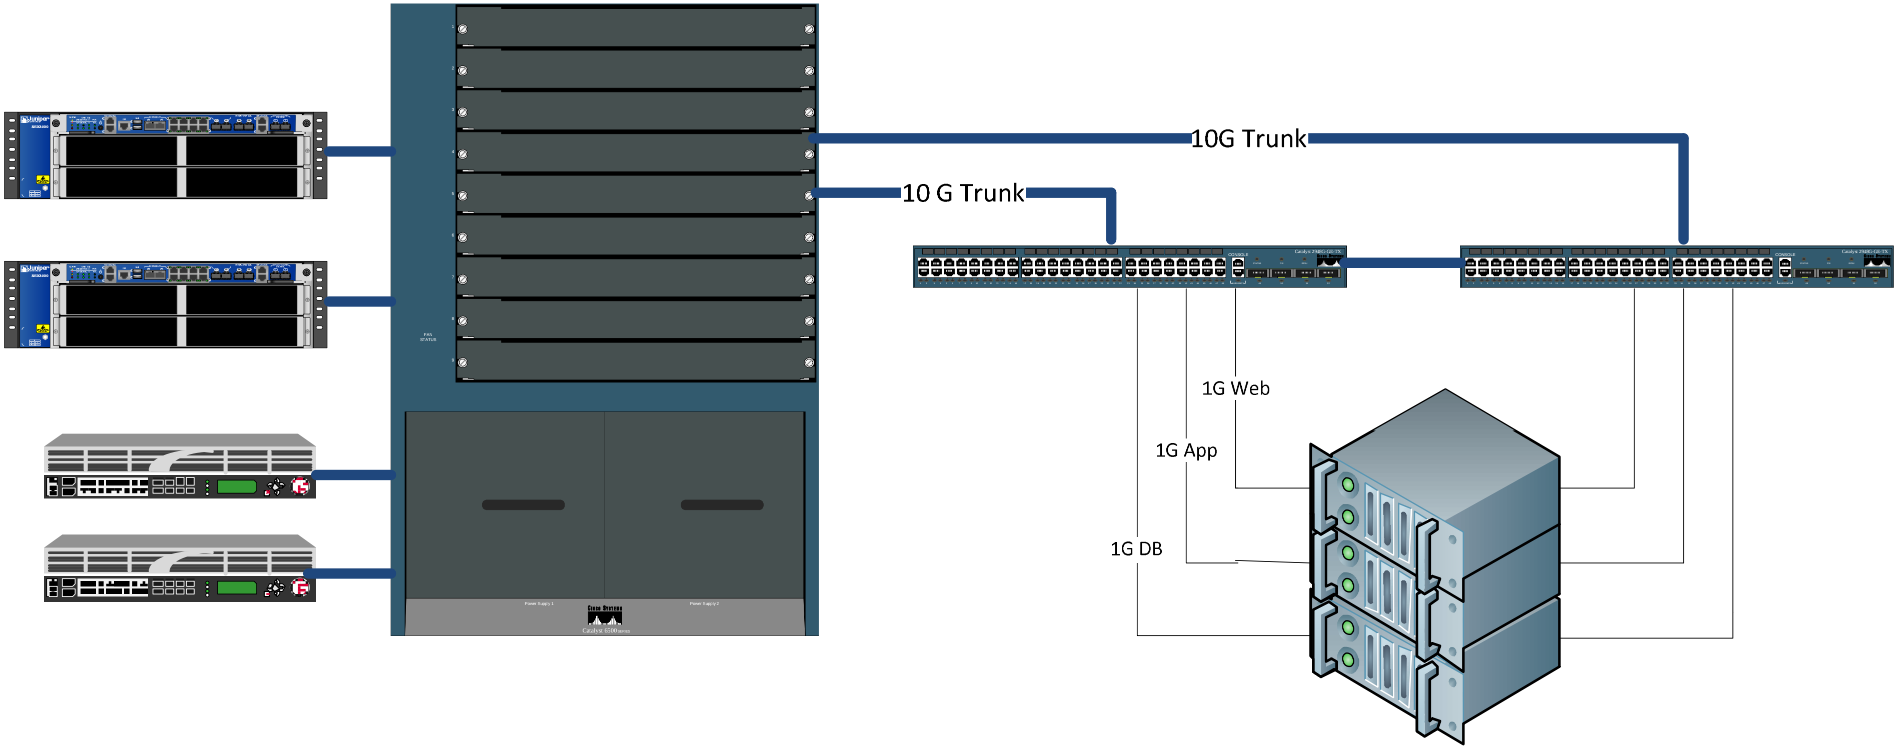

What are we getting from this design?

- Network segmentation
  - Control traffic
    - SSH
    - SNMP
    - RDP
    - File services
  - Production traffic
    - Internal app server ports
    - DB connections

- Network configuration simplicity
  - Slicing up internal VLANs
  - No on-machine firewall, just IP config

What is enforcing our segmentation?

- Custom OS on Routers
  - ACLs
  - Routes

- Custom OS no Switches
  - VLANs
  - Port Security
  - L3 ACLs?

- Custom OS on Firewalls
  - Stateful rules
  - VPN
  - IDP?

Does this still make sense?

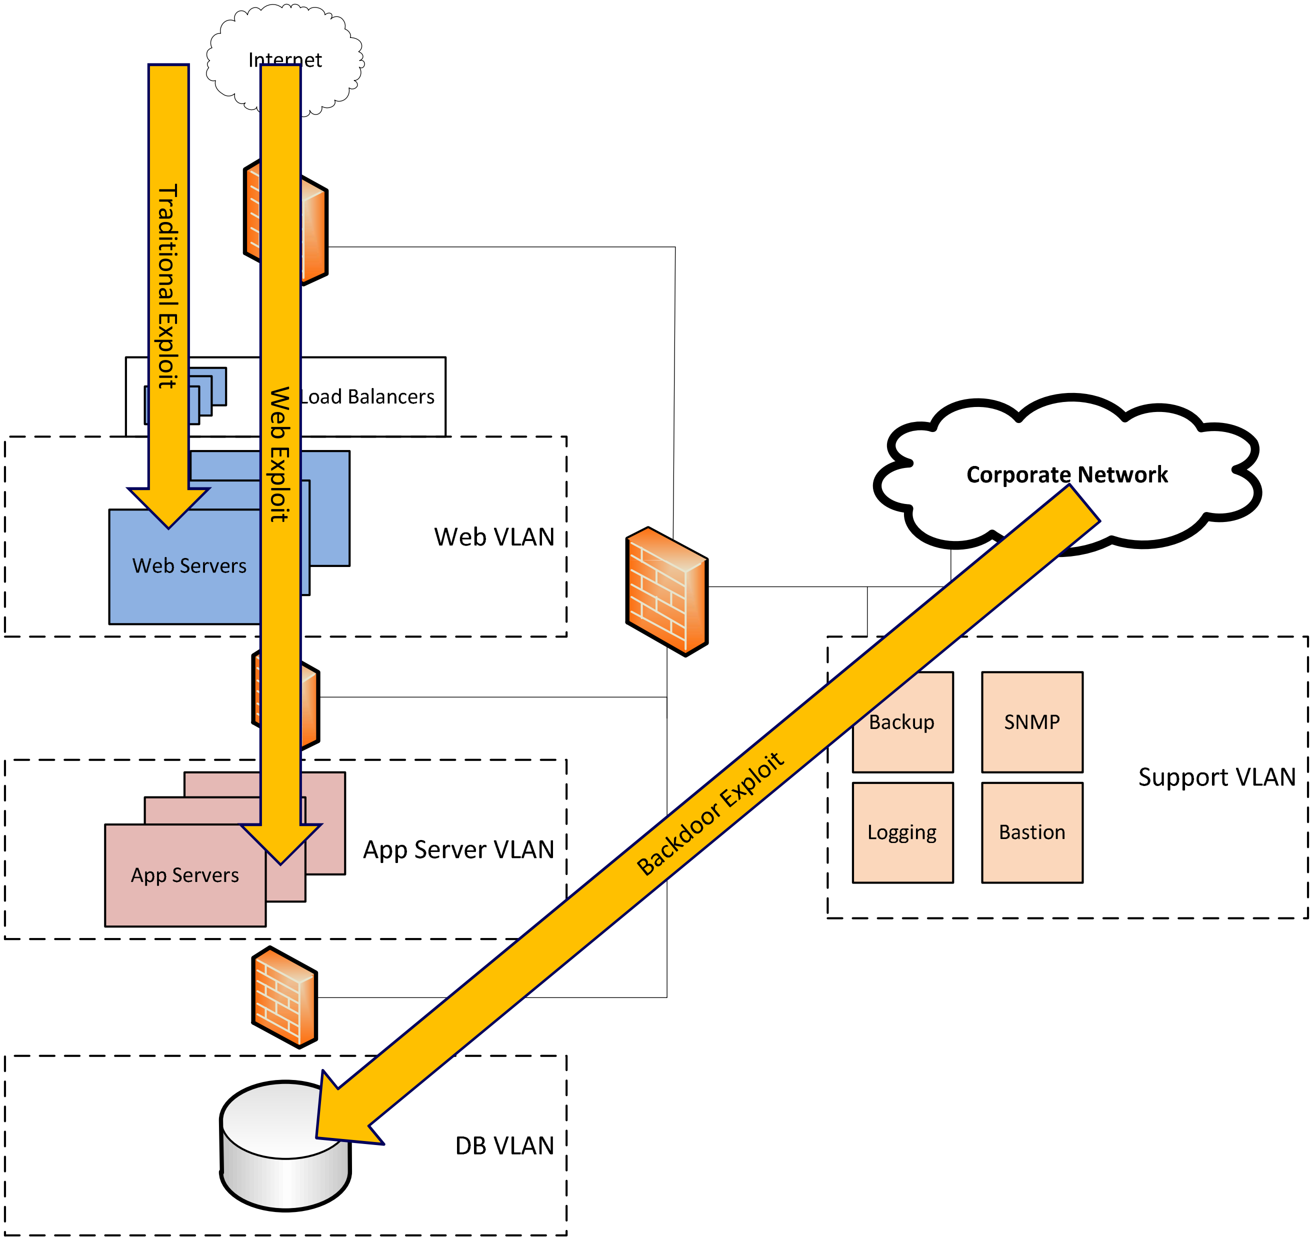

“Traditional Cloud Architecture”

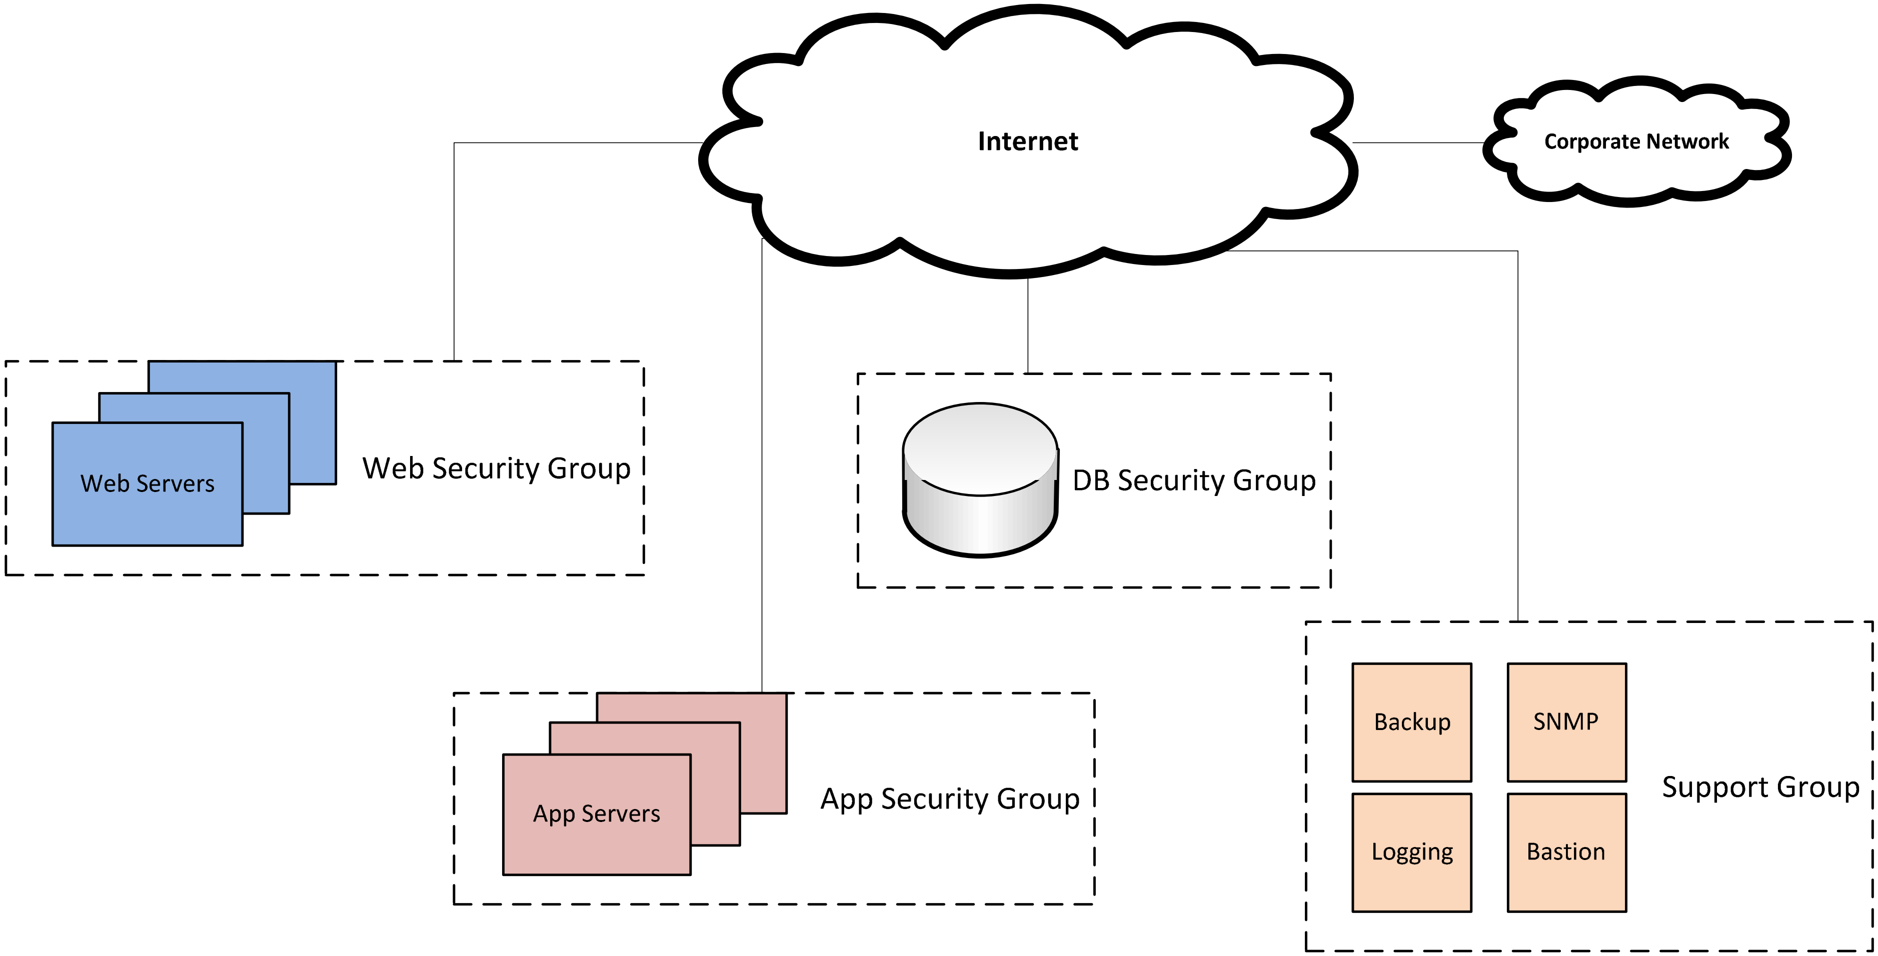

Problems with the "Traditional Cloud"

- Classical network segmentation now harder
- Per-tier scalability is problematic
- Relies on manual creation of relationships
  - How does the proxy server find the app server?
  - How does the app server authenticate to the DB?
- DB performance will be pretty bad

What is the alternative?

- Go Flat
- Collapse the Perimeter
- Use cloud glue services
- Enforce access control via cryptography

Go Flat

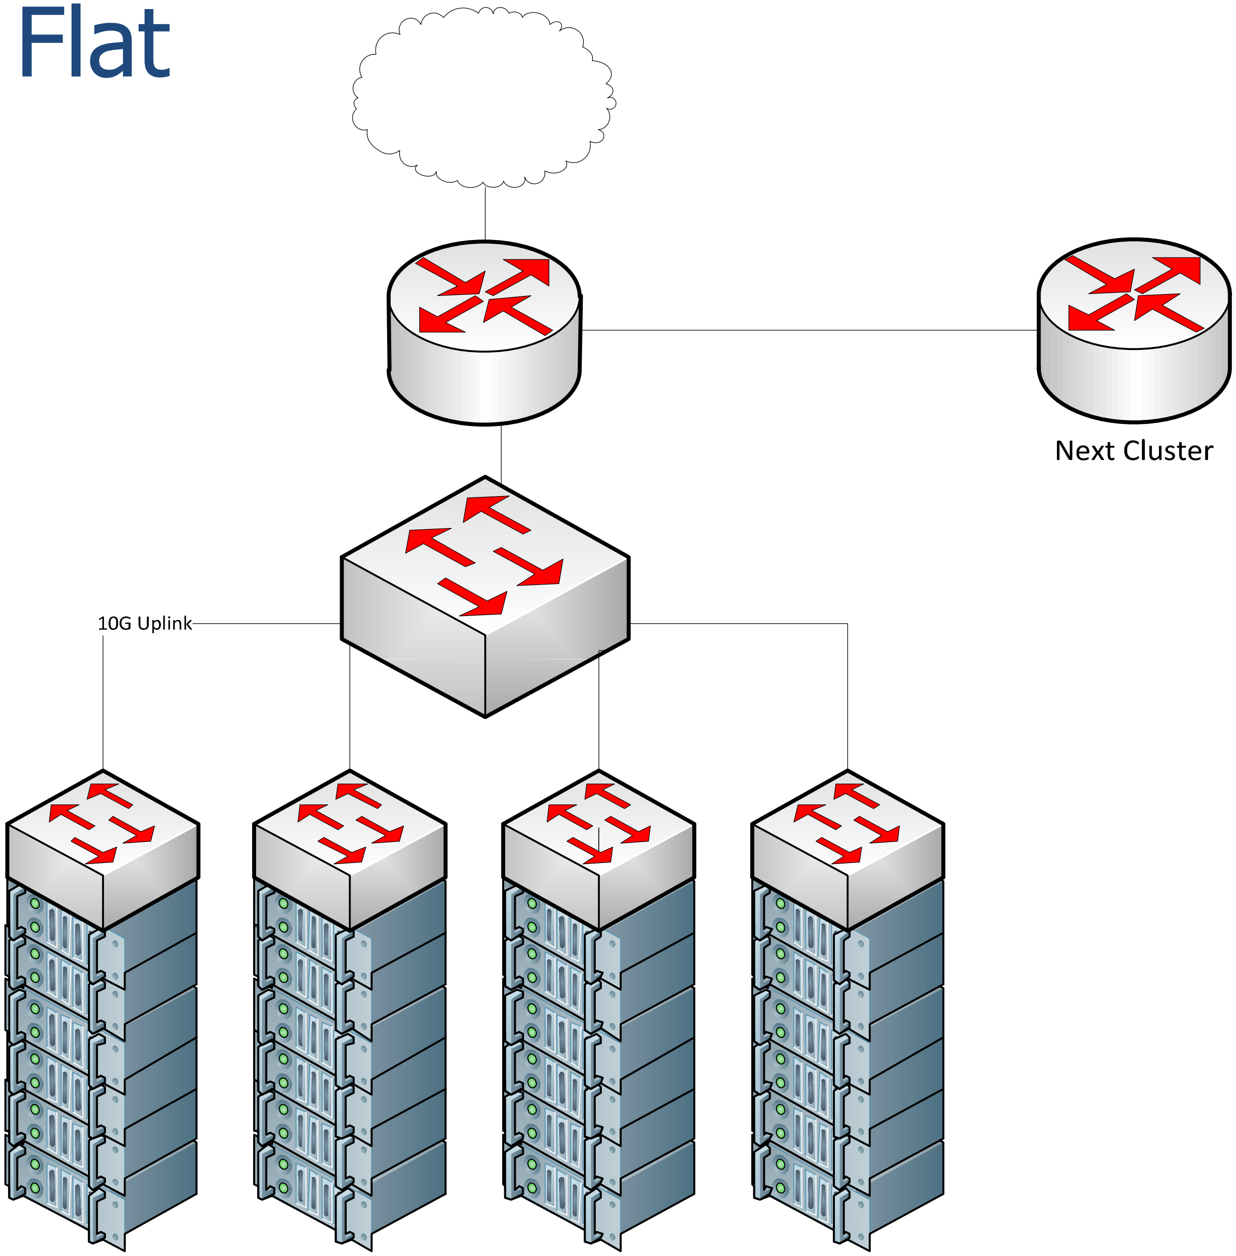

iSEC PARTNERS

Why Go Flat?

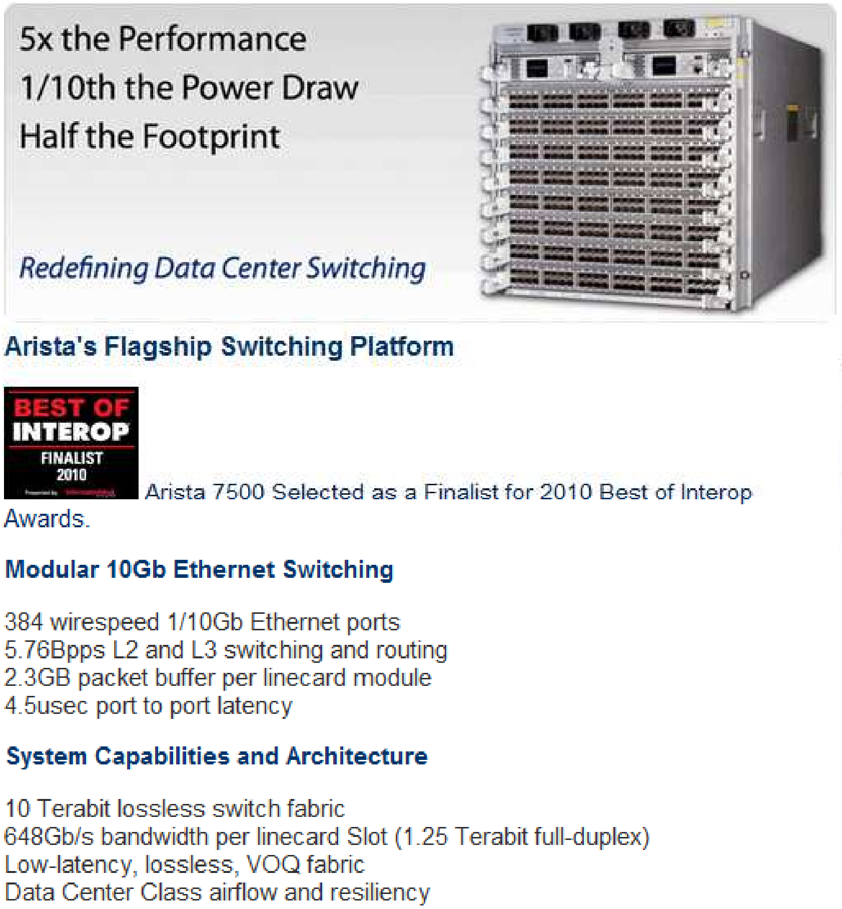

5x the Performance
1/10th the Power Draw
Half the Footprint

Redefining Data Center Switching

Arista's Flagship Switching Platform

BEST OF INTEROP FINALIST 2010 Arista 7500 Selected as a Finalist for 2010 Best of Interop Awards.

Modular 10Gb Ethernet Switching

384 wirespeed 1/10Gb Ethernet ports
5.76Bpps L2 and L3 switching and routing
2.3GB packet buffer per linecard module
4.5usec port to port latency

System Capabilities and Architecture

10 Terabit lossless switch fabric
648Gb/s bandwidth per linecard Slot (1.25 Terabit full-duplex)
Low-latency, lossless, VOQ fabric
Data Center Class airflow and resiliency

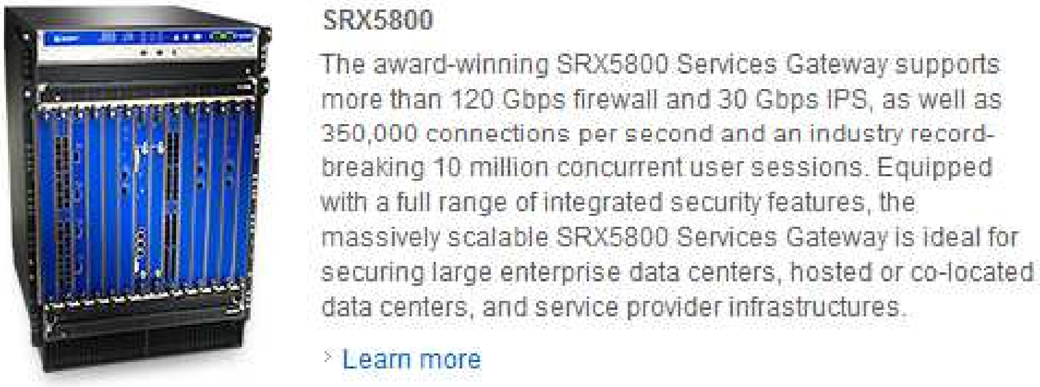

SRX5800

The award-winning SRX5800 Services Gateway supports more than 120 Gbps firewall and 30 Gbps IPS, as well as 350,000 connections per second and an industry record-breaking 10 million concurrent user sessions. Equipped with a full range of integrated security features, the massively scalable SRX5800 Services Gateway is ideal for securing large enterprise data centers, hosted or co-located data centers, and service provider infrastructures.

› Learn more

iSEC PARTNERS

Why Go Flat?

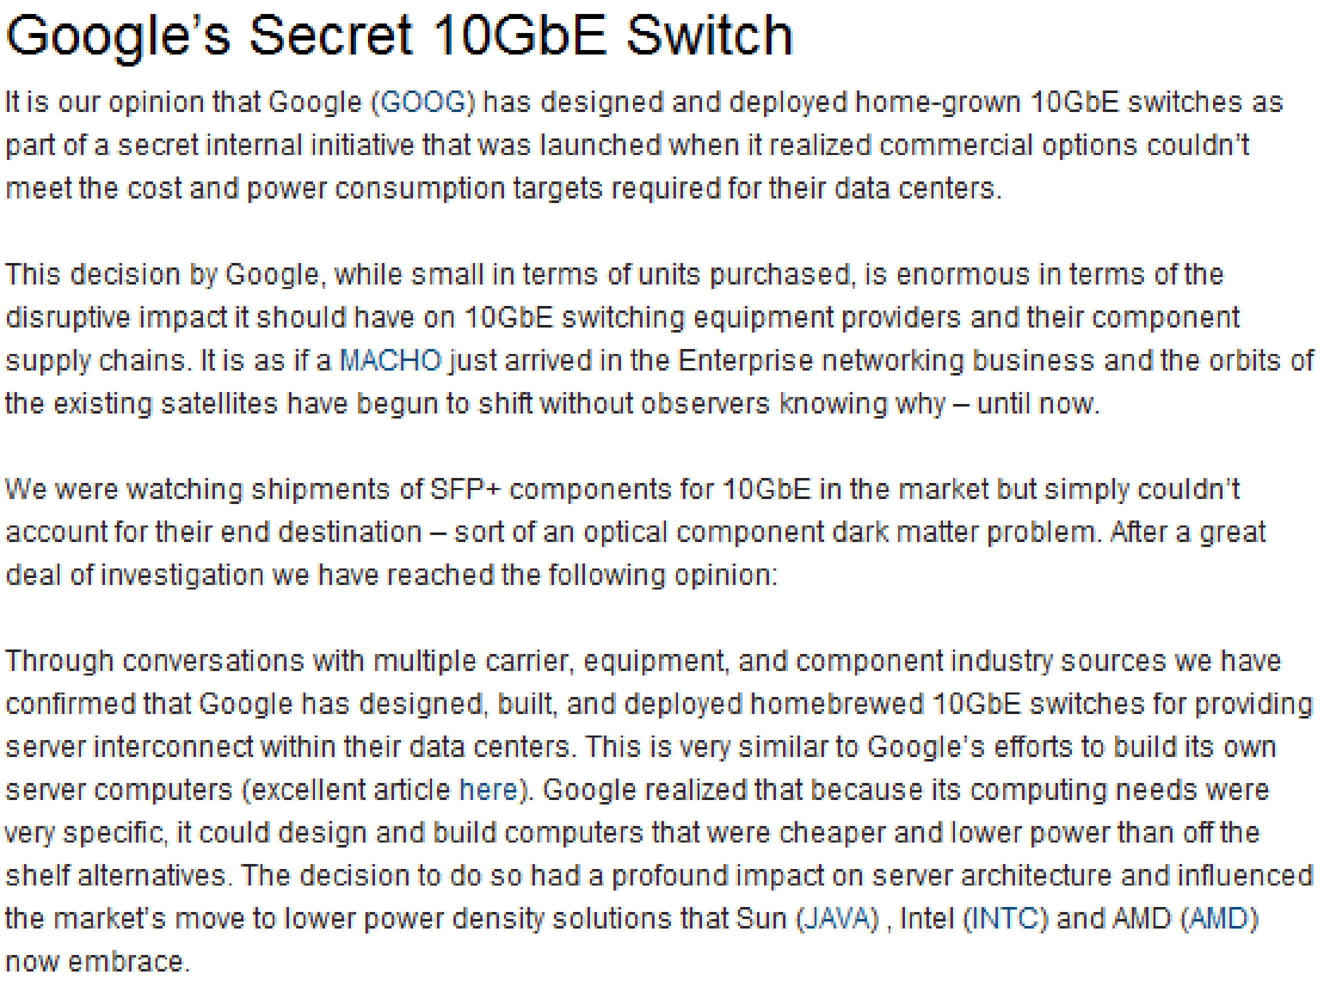

Google's Secret 10GbE Switch

It is our opinion that Google (GOOG) has designed and deployed home-grown 10GbE switches as part of a secret internal initiative that was launched when it realized commercial options couldn't meet the cost and power consumption targets required for their data centers.

This decision by Google, while small in terms of units purchased, is enormous in terms of the disruptive impact it should have on 10GbE switching equipment providers and their component supply chains. It is as if a MACHO just arrived in the Enterprise networking business and the orbits of the existing satellites have begun to shift without observers knowing why – until now.

We were watching shipments of SFP+ components for 10GbE in the market but simply couldn't account for their end destination – sort of an optical component dark matter problem. After a great deal of investigation we have reached the following opinion:

Through conversations with multiple carrier, equipment, and component industry sources we have confirmed that Google has designed, built, and deployed homebrewed 10GbE switches for providing server interconnect within their data centers. This is very similar to Google's efforts to build its own server computers (excellent article here). Google realized that because its computing needs were very specific, it could design and build computers that were cheaper and lower power than off the shelf alternatives. The decision to do so had a profound impact on server architecture and influenced the market's move to lower power density solutions that Sun (JAVA) , Intel (INTC) and AMD (AMD) now embrace.

http://www.nyquistcapital.com/2007/11/16/googles-secret-10gbe-switch/

Collapse the Perimeter

- Per-OS or Per-Hypervisor Software Firewall
- Software load balancer with TLS termination
- Use secure control protocols for top-down management
- Per-device PKI

Can I Trust a Software Firewall?

- Worried about outside access, use routing to segment internal and external systems
  - Can use separate broadcast domain/physical network for public IPs
- Cannot trust for egress filtering. Use static routes to limit access for private IPs to proxy servers
- Can still have non-blocking IDS, although speed is a problem
- Need to treat every machine as a secure individual unit

Overlay Networks

- Another option if you don’t trust your network: overlay IPsec

- Easy with OpenBSD, KAME tools on Linux

- Need per-host certificate
  - No easy open-source option right now

Use Cloud Glue Services

- Variable scalability means loosely coupled applications

- Can use:
  - Asynchronous web service calls
  - Message bus
  - Cloud provided Queuing Service (like SQS)

Loosely Coupled

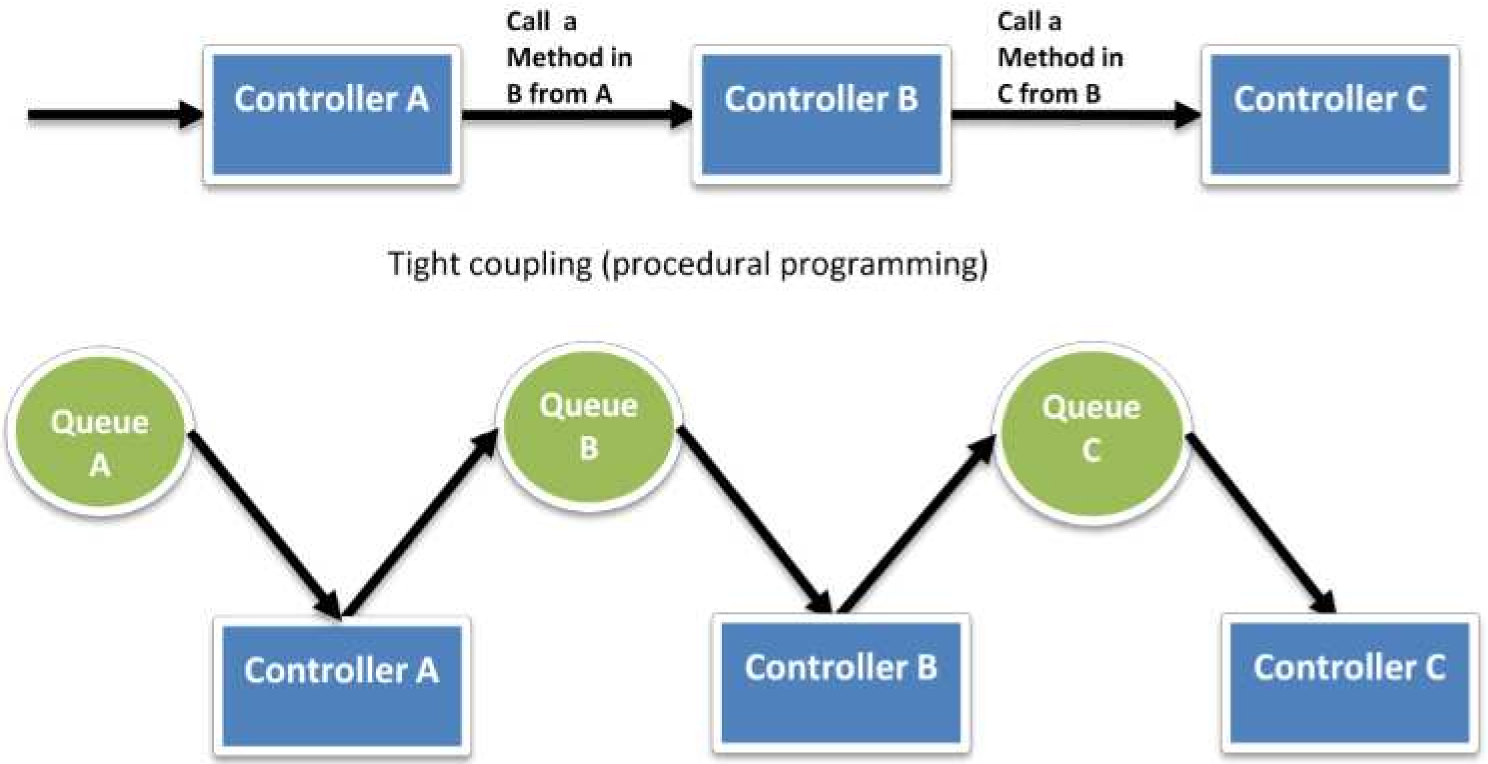

Tight coupling (procedural programming)

Loose coupling (independent phases using queues)

http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

Database Replacement

- Running single-instance MySQL is bound to disappoint
- What are your real data consistency requirements?
  - If you use memcached already, then answer is “none”
- Lots of great work on NoSQL DBs
  - Document Store

    CouchDB relax   Jackrabbit
  - Name/Value

    MEMCACHED   BigTable   SimpleDB

Can You Trust These Services?

- Maybe, maybe not…

But you don't have to.

- End-to-end encryption with symmetric keys between endpoints

Access Control via Cryptography

- With NoSQL, you already are losing a lot of cross-table capabilities
  - Might as well take the plunge…
- Consider per-entity encryption
  - Generate key on entity creation
    - User
    - Institution
  - Store key in "authentication server"
    - Web service request on login
    - Runs in separate cloud or administrative domain
  - Encrypt private fields in the application
    - Helps with cloud and web app security issues

Centralized Configuration Management

- How do you safely get secrets onto cloud VMs?
  - AMIs can’t be encrypted
  - Manual provisioning misses the point

- One option, chroot from shared storage
  1. Boot script on instance creation with built-in key
  2. Download opt.tar.gz.gpg and etc.tar.gz.gpg from S3 bucket
  3. Decrypts using symmetric passphase gpg
  4. Change into chroot
  5. Symmetric key is gone when script is complete

So where did we end up?

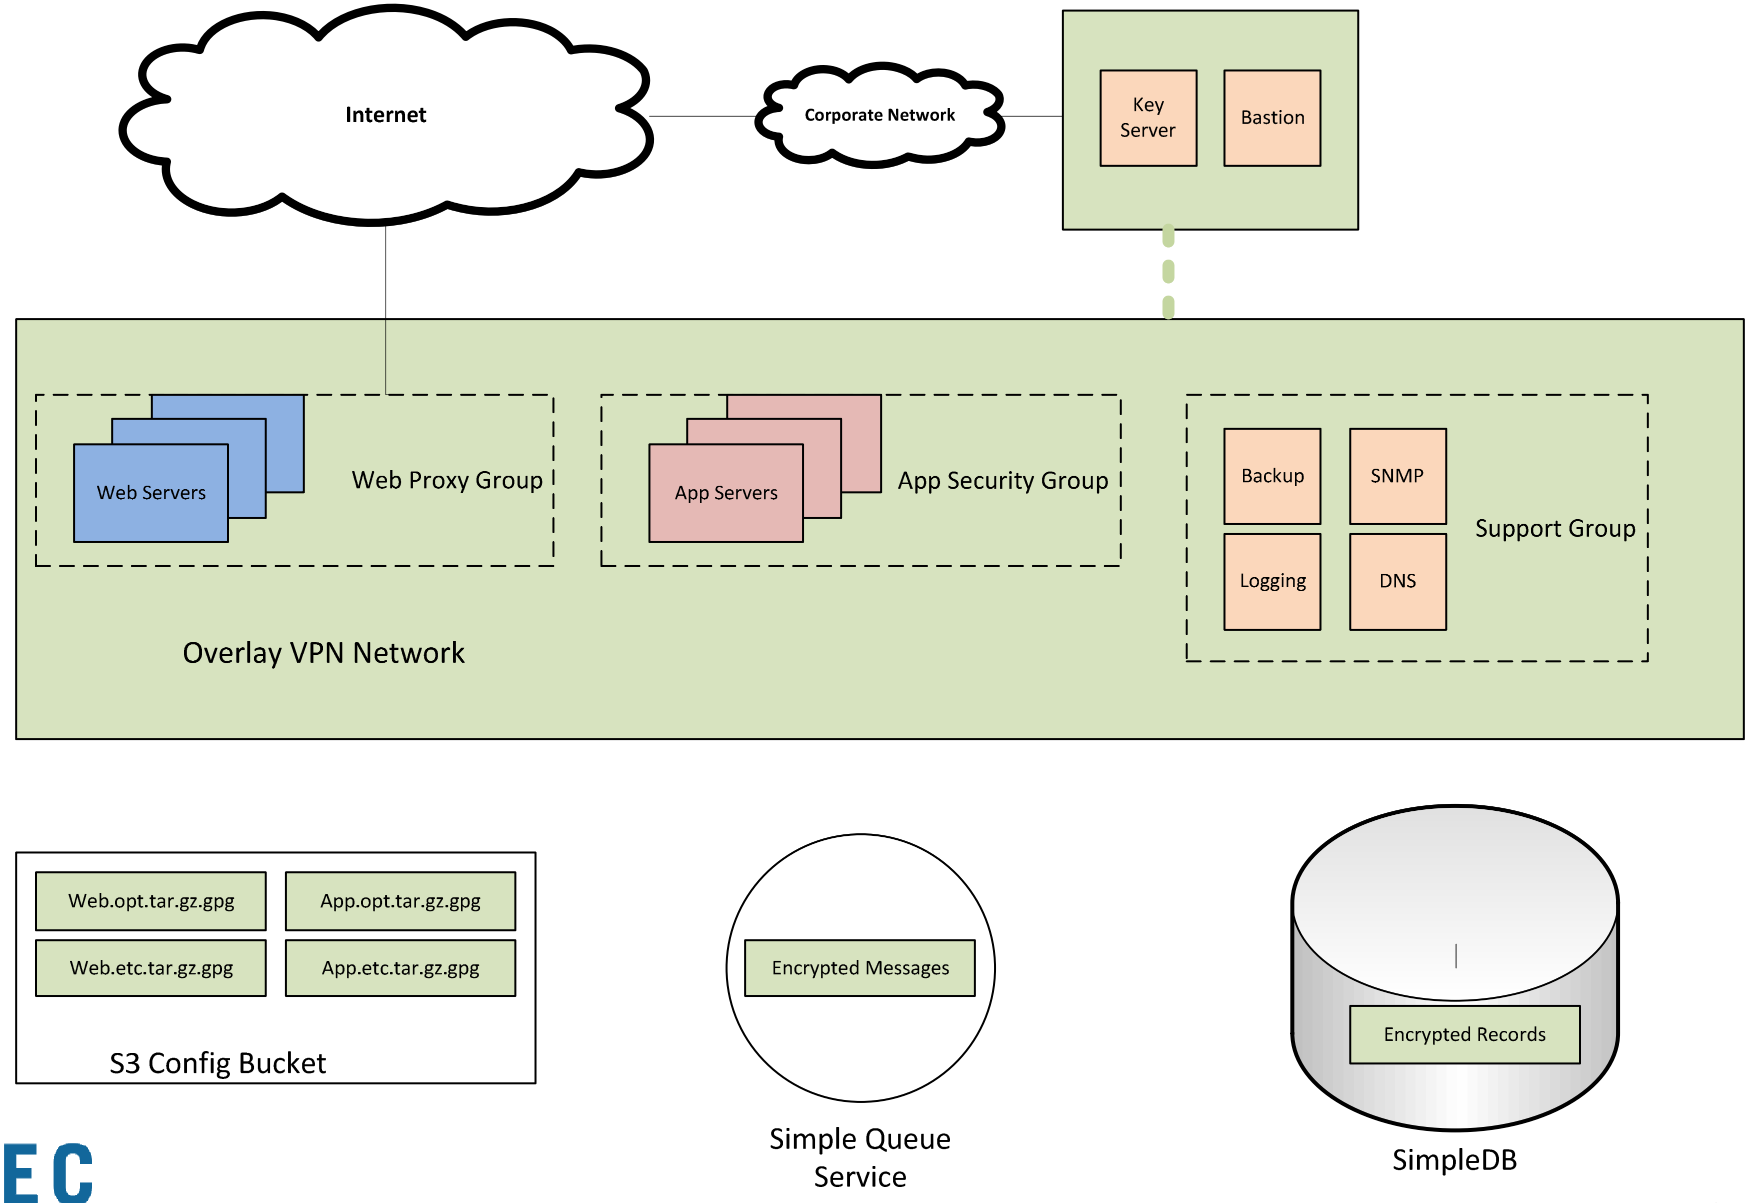

Internet

Corporate Network

Key Server

Bastion

Web Servers

Web Proxy Group

App Servers

App Security Group

Backup

SNMP

Logging

DNS

Support Group

Overlay VPN Network

Web.opt.tar.gz.gpg

App.opt.tar.gz.gpg

Web.etc.tar.gz.gpg

App.etc.tar.gz.gpg

S3 Config Bucket

Encrypted Messages

Simple Queue Service

Encrypted Records

SimpleDB

Still to be done

- Easy PKI
  - Need to:
    - Associate new instances
    - Grant individual asymmetric identities
    - Register in DNS
    - Authenticate requests to key server
    - Receive per-user or per-institution keys
    - Automagically negotiate IKE with peers
    - Revoke dead instances
  - Windows does this much better with AD/IPSec

- Easier configuration management
  - RightScale is moving in the right direction

Still to be done

- Standard PaaS middleware for Enterprises
    - Watch Azure/.Net relationship

- Better "cloud in the box" products
    - Citrix and VMWare are in a horse race

- Adjustment of standards regimes to understand these security requirements

Thank you for coming!

alex@isecpartners.com

iSEC PARTNERS