

Transforming the Information Infrastructure: Build, Manage, Optimize.

FALL 2011

Ethernet – The Basis for SAN / LAN Convergence in the Data Center

SNIA

Greg Scherer

VP, Server and Storage Strategy

gscherer@broadcom.com

Broadcom Corporation





Agenda

- Broadcom Introduction
- Reasons for Convergence in the Data Center
- Key Market Drivers for Convergence
- Technology Enablers for Convergence
- Adapter Capability and Architectural Requirements
- Conclusions and Call to Action





About Broadcom

- Global leader in semiconductors for wired and wireless communications
- Founded 1991
- Fortune 500 company
 - Ranked most innovative semiconductor
- 2010 Gartner Top 10 Semiconductor Companies (Revenue)
 - 2010 net revenue of \$6.82 billion
- One of the largest volume fabless semiconductor suppliers
- Broad IP portfolio with over 15,000 U.S. and foreign patents and applications
 - Strongest patent portfolio among fabless semiconductor companies (IEEE)
- Approximately 9,052 employees worldwide

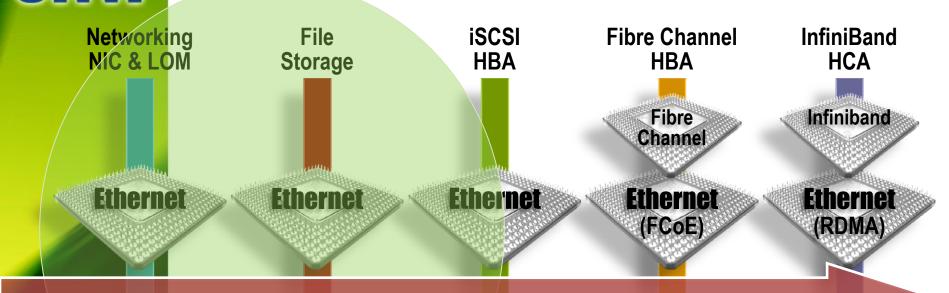






Driving convergence through unique & innovative product offerings

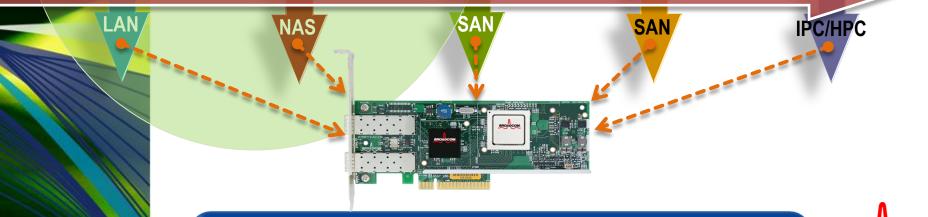
Ethernet: Consolidation Followed by Convergence (Multiple Networks vs. one Unified Network)



SNIA

COMPUTERWORLD

Convergence Trend Line Over Time



Consolidation over Ethernet is underway; Convergence is Next!





Ethernet:

_eading

Economics

Networking and Storage Convergence *1. Ethernet* Why?

Open, Interoperable, Simple, Low Cost, and High Volume!

2. Unified & Flat Networks

Significant Opportunity for Lowering CAPEX and OPEX

3. Virtualization

Network Storage is Required for Migration, DR, etc.

4. Scalable Bandwidth

Clear Path Beyond 10Gb to 40Gb, 100Gb, etc.

5. Continued Innovation

Loss-Less (DCB), Energy Efficient Ethernet (EEE), Congestion Mgmt (802.1Qau), Advanced Routing (TRILL), etc.

Ethernet : Technology leadership fueled by economics



2/4/8G FC 16/32G FC 64G FC

40GbE

ETS

100GbE

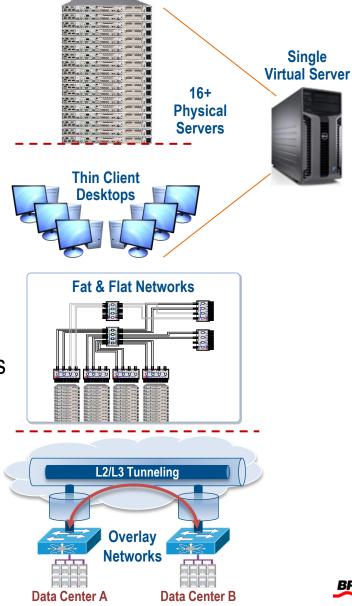
QCN

10GbE



Key Trends Driving Convergence

- Server Virtualization
 - Server consolidation
 - Desktop virtualization



everything^{*}

Network Virtualization

- Fat & Flat L2/L3 networks
 - Driven by public Cloud
- Overlay networks



Cloud Computing Characterization

"Cloud" – a metaphor for on-demand IT <u>"In the Network"</u> <u>Not a technology</u> but an <u>architectural practice</u> Implemented with a suite of technologies <u>Distributed computing environment of arbitrary scale</u> Where hardware and software are abstracted from services



ESSENTIAL CLOUD ATTRIBUTES Leading Economics:

- → Extreme efficiencies do more with less
- → Utility metering pay as you go
- Utilization capacity/demand impendence
- → Business service entry cost

Flexibility:

- \rightarrow Any server can be used for any application
- → Application time-to-market
- No lock in Standard interfaces
- Programmable

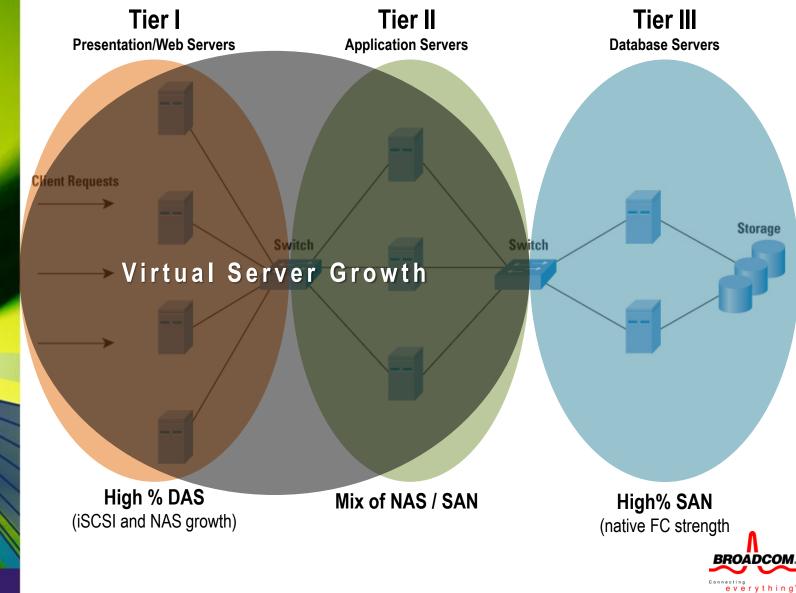
Elasticity:

- Dynamic resource pooling Grow & Shrink
- Migration
- Parking



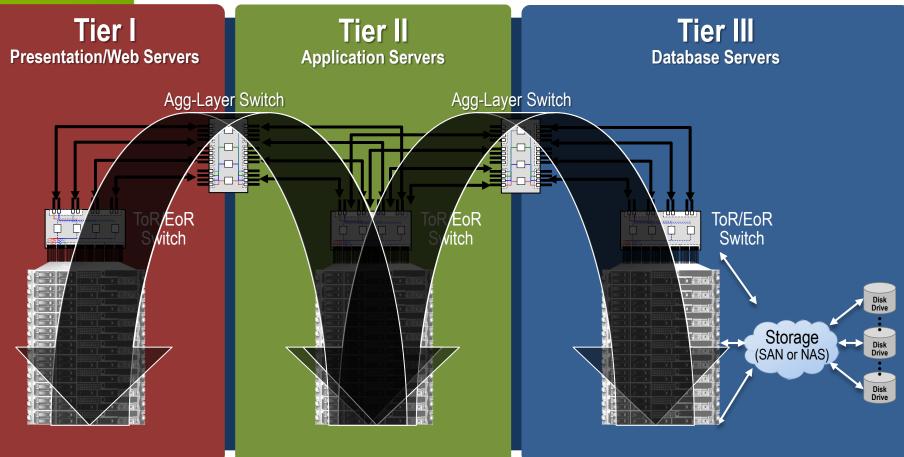


Traditional 3-Tier Enterprise Datacenter Logical View





Traditional 3-Tier Enterprise Data Center Physical View





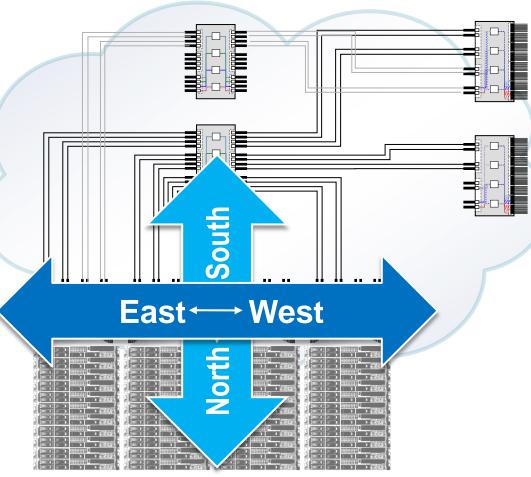
Traffic was predominantly North→South – almost no traffic within each tier Provisioning is a "Physical" exercise and difficult to re-configure





The New "Virtual Data Center" Flat network with full cross-sectional B/W from any server to any server

Fully Connected, Non-Blocking Fabric Interconnect



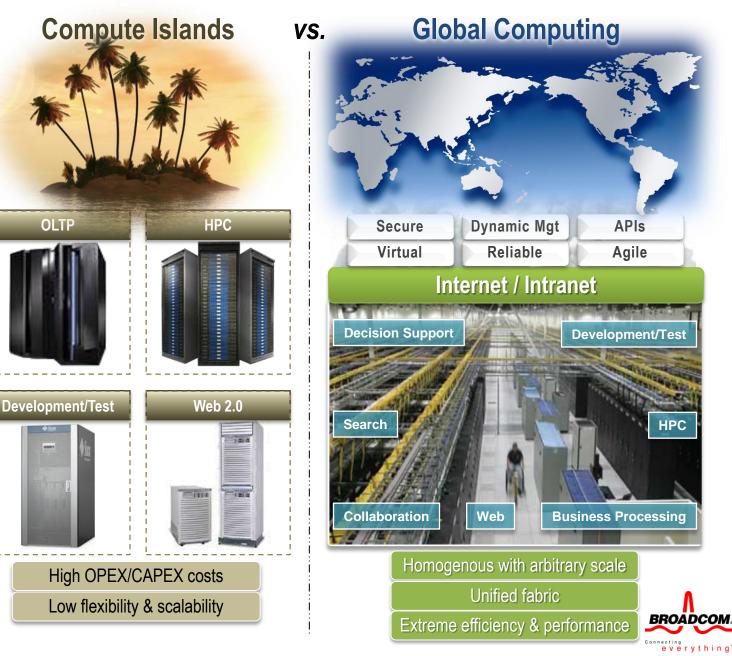
- Any Server can be in any "Tier" of the Data Center
 - Provisioning is completely Virtual
- Fast re-configuration is possible with minimal manpower
- No assumptions can be made regarding North -> South vs. East -> West

traffic patterns





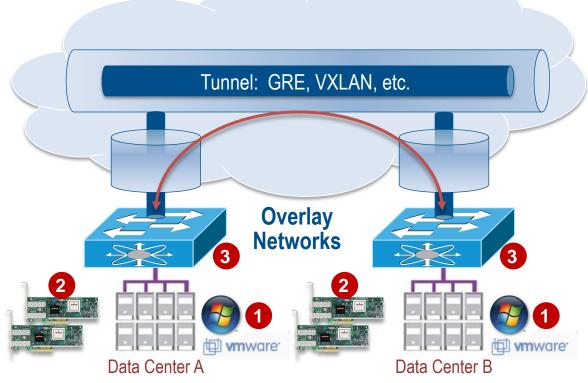
Traditional Approach vs. Public Cloud





Convergence and Overlay Networks (also known as Network Tunnels)

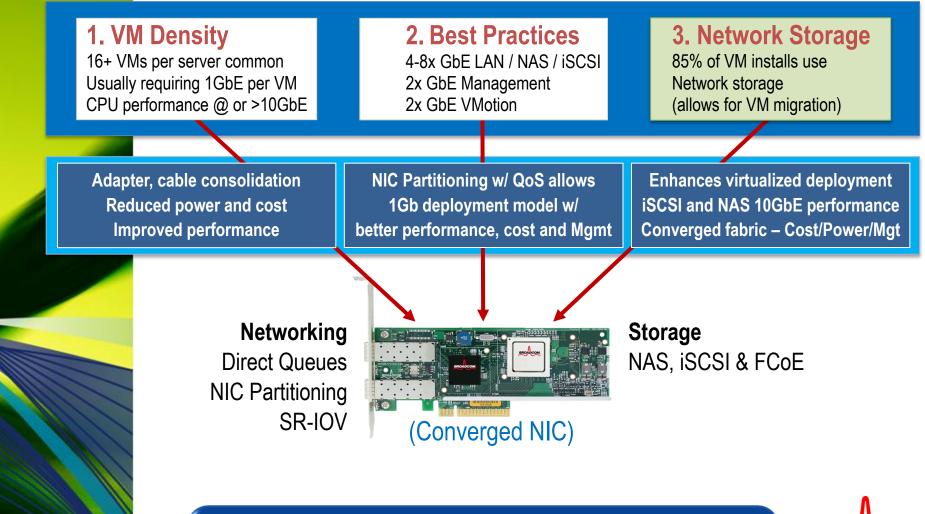
- Tunnels form "Virtual Networks" that fully abstract the Physical Network, extending existing constructs (VLANs, IP subnets, etc.)
- Tunnels can be originated or terminated in:
 - Host software
 - Network controller
 - Fabric switch







Server Virtualization is Driving 10G Vectors Driving Convergence



Server virtualization is a huge driver for 10Gb+ Adoption

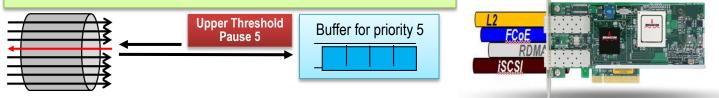
everythin a'



Technologies that Facilitate Convergence PFC, ETS, QCN, and DCBx



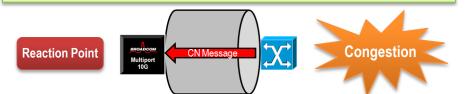
- ✓ Provides link layer flow control by priorities 8 priorities
- ✓ PFC prevents drops during the onset of congestion

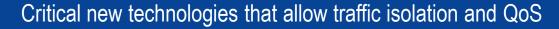


- Enhanced Transmission Selection (ETS): IEEE P802.1Qaz
 - Provides minimum bandwidth guarantees
 - ✓ Unused BW is available for others

5G	3G
4G	4G
	JG (

- 3 Gbps rate limit
- 4 Gbps bandwidth guarantee, High throughput No drop
- 1 Gbps bandwidth guarantee, Best Effort High Throughput
- Quantized Congestion Notification (QCN): IEEE P802.1Qau
 - ✓ Reduce congestion duration by slowing traffic at the src
 - ✓ End nodes monitor buffer utilization for desired equilibrium







"DCBx" Uses LLDP to exchange

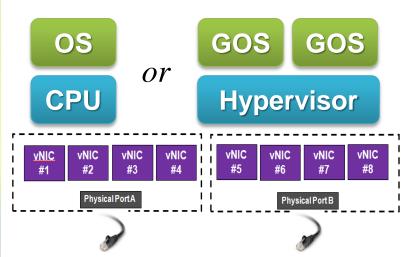
PFC/ETS

configuration

Network Switch

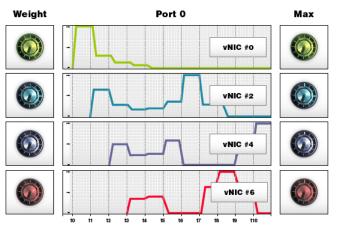


Virtualization: NIC Partitioning



1. Independent NIC partitions

- ✓ Four or more per physical port
- Supports single OS or Hypervisor environments
- ✓ No OS or BIOS changes required
- Switch independent and switch dependent modes of operation



2. Fine-Grain B/W Control

 Control for fixed Bandwidth caps and Bandwidth weighted priority



3. Convergence

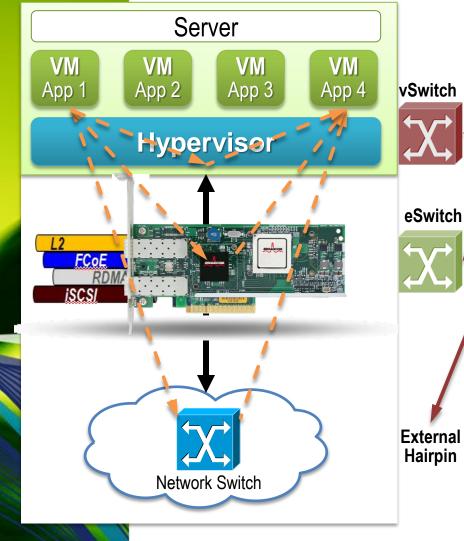
✓ Networking and storage



NIC partitioning: Multiple logical pipes on a single big pipe with QoS



Server Virtualization: Managing the New Switching Edge



New switching points inside the Ecosystem:

NIC Partitioning and SR-IOV Enable Internal Switching between vNICs and vHBAs

Benefits

- Flexibility and availability
- Speed of innovation

Challenges

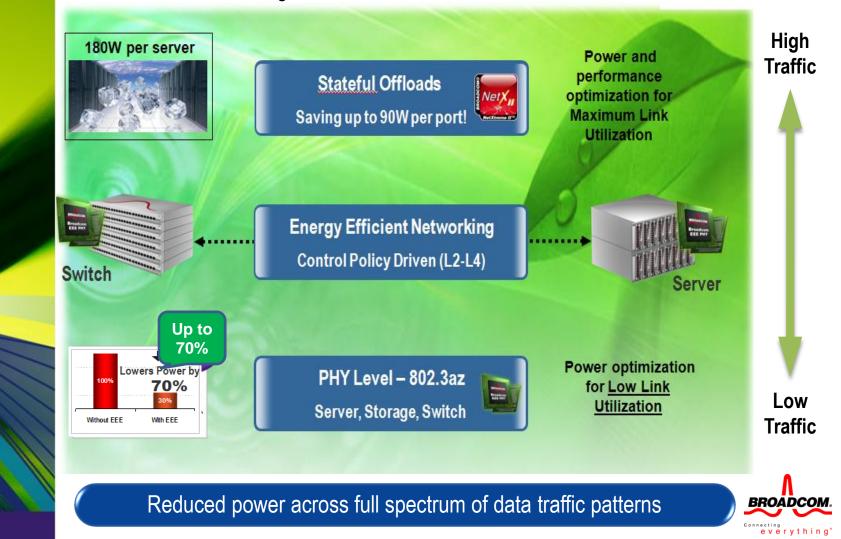
- Management
- Policy enforcement
- Performance (for vSwitch)





Convergence and Power Management Offloads, EEN, and EEE

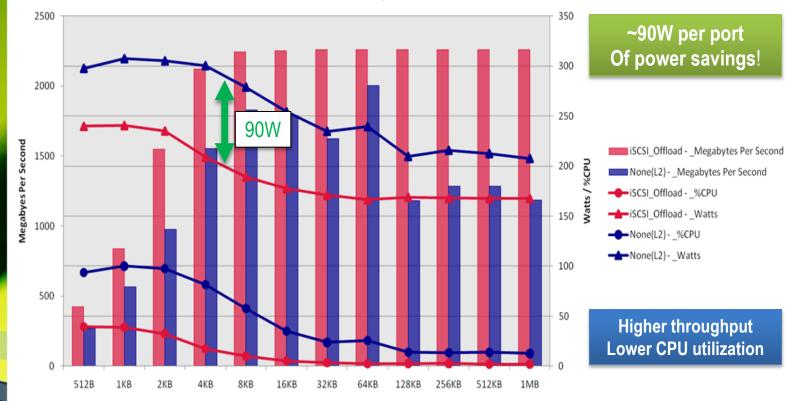
In addition to consolidating Controllers and Switch ports, state-of-the-art Converged Controllers Conserve Power with Offloads!





Power Management in Action iSCSI Hardware Offload Example

Write - MBps v CPU v Watt



Offload Provides Data Centers

- ✓ Improved PUE (Power Utilization Efficiency)
- ✓ Returns CPU to be used for applications, not I/O



COMPUTERWORLD SNIA

Three Tenants of Virtualized Performance



All three elements are required to provide SLAs

1. Handling <u>Consolidated</u> Server Workloads

Increased workload with #of VMs per Server

2. Handling <u>Converged</u> I/O Workloads

- Networking + Storage (NAS, iSCSI or FCoE)

3. Handling *Bursts* of Traffic

- There's no such thing as an "Average" traffic burst
 - When an application *steps on the gas,* it wants immediate response!

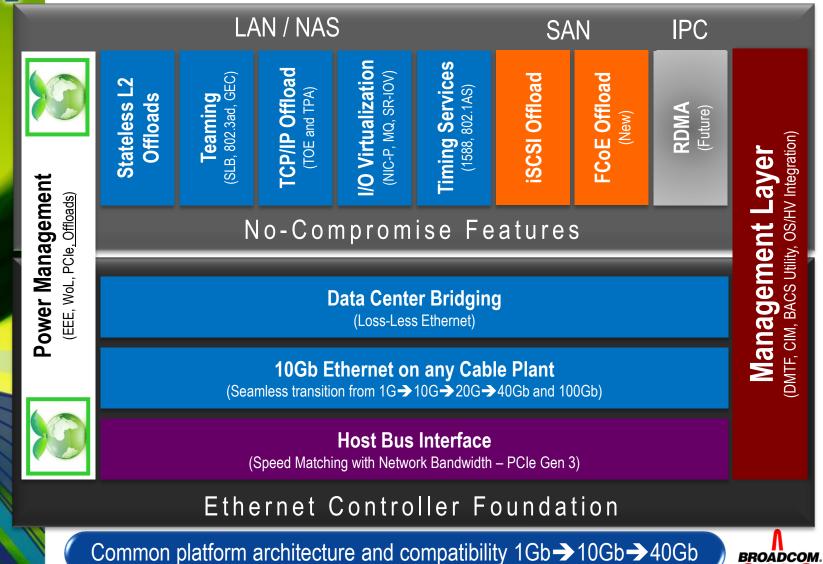




everything



Ideal Converged Controller Architecture

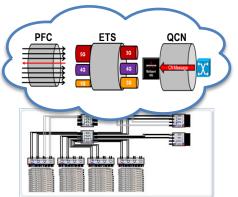


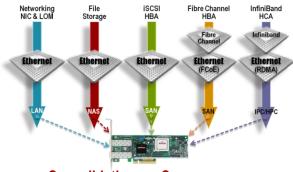




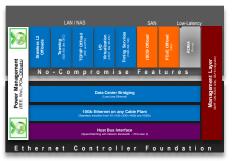
Conclusions and Call to Action

- Consolidation and convergence require new network architectures to be deployed
 - Loss-less Ethernet (DCB)
 - Fat and flat L2/L3 fabrics
- Storage and networking are, and HAVE been converging
 - Server virtualization and Cloud are accelerating the process
 - NAS is already converged with L2 and iSCSI has begun
 - FCoE is just the next workload to migrate to Ethernet
 - Short term consolidation over Ethernet; longer term convergence
- No one knows when full convergence will occur, so...
 - Future-proof with controller and switch architectures that are convergence-ready today!





Consolidation vs. Convergence



everything

ectures today!

Future-proof with convergence-ready architectures today!



Thank You !

Have a question? Contact me at: gscherer@broadcom.com







Useful Links

Hypervisor vSwitch: <u>http://openvswitch.org/papers/hotnets2009.pdf</u>

Multi-Queue: <u>http://www.broadcom.com/collateral/wp/Virtualization-WP100-R.pdf</u>

SR-IOV: http://www.pcisig.com/specifications/iov/single_root/

- EVB: <u>http://wikibon.org/wiki/v/Edge_Virtual_Bridging</u>
- DCB: <u>http://en.wikipedia.org/wiki/Data_center_bridging</u>
- iSCSI: <u>http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf</u> <u>http://www.broadcom.com/collateral/wp/iSCSI-WP200-R.pdf</u>

http://www.broadcom.com/collateral/wp/iSCSI-VMware-WP100-R.pdf

- FCoE: <u>http://fcoe.com/</u>
- EEE: <u>http://www.itu.int/dms_pub/itu-t/oth/09/05/T09050000010005PDFE.pdf</u>





Transforming the Information Infrastructure: Build, Manage, Optimize.

FALL 2011