Gandalf Overview

Steven Wood

Contributors

Gunes Aybay	– Riverstone Networks
Mark Bordogna	– Agere Systems
David Cheon	– Sun Microsystems
Permider Chohan	– Infineon Technologies
Spencer Dawkins	– Fujitsu Network Co.
Martin Green	– Cisco Systems
Yong Kim	– Broadcom
Sateesh Kumar	– Redwave Networks
Dave Meyer	– Mindspeed
Gal Mor	– Corrigent Systems
Chuck Lee	– Appian Communications
Ashwin Moranganti	– Appian Communications
David James	– Lara Networks
Bob Sultan	– DataNet Associates

Gandalf Overview

- This presentation provides an overview of the MAC features contained in Gandalf
- 13 different organizations participated through conference calls and comment resolution methods on the draft
- Gandalf proposes a number of compromises on some of the more contentious issues facing 802.17
  - Deeply held beliefs on both sides require a certain amount of dual mode operation in the standard for progress to be made
  - Allows people to optimize cost for performance

Gandalf Overview

- Gandalf is not compatible with SRP
  - New features and capabilities added
  - SRP nodes and Gandalf nodes do not inter-work
- All silicon must be respun
  - Changes provide greater capabilities while allowing standard product 802.17 silicon to be available to the entire market rapidly

Frame Format

- Identification of a *steering only* data packet
- New ring control bit to mark In/Out_of_Profile packets
  - supports a medium priority service
- Addition of a HEC field in the frame protecting from the start of frame through the Type field
  - allows delivery of TDM packets with payload errors
- ***Not backwards compatible with any existing implementation***

Passthru Path Design

- Base case for the standard is either one of
  - Dual Passthru buffers – store and forward (SF) operation
  - Single Passthru buffer – SF or cut through (CT) operation
    - PB is the same size as the HP buffer in the dual case
- Fairness Algorithm operates well
  - Dual PB design at 100% utilization
  - Single PB design at 90-95% utilization
  - In mixed rings single PB nodes are not disadvantaged
  - Scalable to a large number of nodes and ring distance
- Advantages of single buffer
  - Slightly reduced complexity in transmit path
  - Slight reduction in cost
    - Can be ignored in the overall cost of a system

Passthru Path Design

- Advantages of dual buffer
  - better high priority jitter and latency
    - low priority buffer virtually eliminates blockage of HP transmit traffic
  - ability to fully utilize all link BW
    - no requirement to limit utilization to 90%
  - simple node based fairness algorithm
    - Small number of state variables = stability.
    - Minimal BW message Tax
  - Congestion-Avoidance Avoidance
    - low priority PB able to absorb burst behavior – removes the need for congestion avoidance
  - cut through operation not required except at very low speeds

Fairness Algorithm Overview

- Node based fairness
  - Nodes operate independently – no need to communicate with all other nodes on the ring
  - Tolerant to lost messages
    - Nodes slowly speed up when no rate information arrives
    - Fairness messages also used as keep-alives
- Fairness messages propagate backwards on opposite ring
  - Congestion domain isolation
    - Message only travels as far as necessary
    - Stops when it reaches a node that is conforming to the advertised rate
    - BW used to communicate rate information minimized
    - interactions between nodes minimized

Fairness Algorithm

- Simulations have shown it operates well with both single and dual PB designs
  - allows homogeneous systems to operate well
    - cost “difference” is an implementation choice
  - allows interoperation of single and dual systems
- Simulation comparison to Alladin proposal are not complete yet
  - we simulated a simple scenario on the Alladin model provided
  - as a courtesy the results were provided to Alladin group for comment
  - Alladin group informed us there was a problem with the current simulator. New one released.

Gandalf is not SRP

- Fairness method has many extensions not supported by SRP
  - None of these features are implemented by silicon in design or available today
  - Pre-provisioned High Priority traffic
    - Limit fairness to X% of ring BW
  - Weighted Fairness
  - Multi-Choke Point operation for VOQ support
  - Medium priority profiler / traffic marker
    - Behavior is different for in and out of profile traffic
  - Operation with both single and dual buffers

Protection

- Gandalf proposes the standard supports either wrapping or steering on homogeneous rings
  - heterogeneous rings support wrapping (SWIS)
    - packets can be designated steer only
- Wrapping provides fastest response to failure
  - no need to communicate with all other stations for corrective action
  - lowest packet loss
  - no special cases for multicast packets
- Steering does not require a passthru path to the opposite MAC yielding simpler HW
  - not an issue for a dual ringlet MAC chip
- Steering re-optimizes traffic flow
  - steering following wrap can do the same, best of both worlds

Physical Layer

- Media independent MAC
  - reconciliation layers provide MAC <-> PHY translations
- SONET PHYs include
  - POS
  - GFP
    - optional behavior for MAC to provide length to GFP framer
    - many framers support length calculation
    - additional queueing delay not a real concern above OC-12
- Ethernet PHYs include
  - 1 GB and 10 GB
  - IPG, preamble, frame delineation are reconciliation layer functions

Gandalf v. SRP

Technology	SRP	Gandalf	Notes
Dual-buffer Transit Path	✔	✔	
Congestion Management	✔	✔	
Congestion Domain Isolation	✔	✔	
Low Delay/Jitter Class	✔	✔	
Best Effort Class	✔	✔	
Topology Discovery	✔	✔	
Wrap Protection	✔	✔	
Steer Protection	✘	✔	
Guaranteed BW Class	✘	✔	
Single-buffer Transit Path	✘	✔	
Customer Separation	✘	✔	
Ring Header Protection	✘	✔	
VOQ Support	✘	✔	
Weighted Fairness	✘	✔	
POS/Ethernet/GFP PHY	✔	✔	

Gandalf Summary

Technology	Gandalf	Notes
Transit Path Buffering	Single or Dual TB	
Transit Path Staging	Store & Fwd	
Congestion Awareness	Local, Local+Global	
Fairness Messaging	Hop-by-hop	Caveat: multichoke
VOQ Support	Yes	
Weighted Fairness	Yes	
High Priority Add Rules	Open a Slot	
Low Delay/Jitter Class	Yes	Utilization Dependent Access Delay
Guaranteed BW Class	Yes	Implementation Differs
Best Effort Class	Yes	

Gandalf Summary

Technology	Gandalf	Notes
Topology Discovery	Hop-by-hop	
Protection	Wrap & Steer	
Frame Format	Ring Header Based	
Ring Header Protection	Yes	
Customer Separation	Optional	
OAM Support	Common Draft	
Station Synchronization	Slip Buffer	Separate timing domains
POS/Ethernet/GFP	Yes	
802.1D Bridging	Yes	
Enhanced Bridging	Encap/RC	Goals: spatial reuse

Conclusions

- Gandalf
  - meets objectives set by 802.17 WG
  - allows system vendors to trade off various cost points and system design considerations
    - PB, Protection
  - will enable rapid time to market for standard product silicon
- Gandalf is not compatible with SRP