

The background of the image is a collage of various items: a magnifying glass with an Apple logo on its handle, a green pen holder with several pens, a globe, and some papers. The text is overlaid on this background.

Worldwide

Developers

Conference



Advanced Mac Networking & Communication

Thomas Weyer

Network & Communication
Evangelist

Advanced Mac OS Networki and Communications

Session 109, Thurs. 9:50-10:50 am, Rm

- Introduction and Logistics
- Session Overview
- Other Sessions of Interest
- Contact Information



Session Overview

- Introduction
 - Thomas Weyer (weyer@apple.com)
- Single Link multihoming
 - David Kim (davidkim@apple.com)
 - Chin Niazi (cniazi@apple.com)
- OT 1.1.x API Changes
 - John Comiskey (comisk.j@apple.com)
- Performance Optimization
 - Quinn (eskimo1@apple.com)
- Q&A
 - DTS & OT Engineering



Previous Sessions of Interest

Mac OS Networking & Communications

- Session 108, Wed. 5:50-6:50 pm, F
- Agenda:
 - Evangelism Introduction
 - Thomas Weyer (weyer@apple.com)
 - OT Roadmap and update
 - Richard Ford (rvf@apple.com)
 - ARA & OT PPP Roadmap and update
 - Richard Ford (rvf@apple.com)
 - Q & A
 - DTS & OT Engineering



Related Sessions

- Rhapsody Core OS Kernel and Runtime
 - Session 201, Mon—4:50-5:50 pm, H
- Mac OS Support in Rhapsody's Blue
 - Session 204, Wed—1:30-2:50 pm, H
- Rhapsody Network Administration M
 - Session 221, Fri—4:30-5:30 pm, Roc
- Rhapsody Networking APIs & Servic
 - Session 220, Fri—5:50-6:50 pm, Roc



Related Sessions

- Feedback Forums
- Rhapsody Core OS Feedback Forum
 - Session 292, Thurs.—11:10-12:10 p.m.,
Room J-4
- Mac OS Networking Feedback Forum
 - Session 194, Thurs.—1:50-2:50 p.m.,



Important Contact Information

- ARA 3.0
 - Blanket NDA Seeding
 - CCL submission deadline of June 1st (msg@apple.com)
- Third-party OT mailing list
 - <http://www.stairways.com/maillinglists/opentpt.html>
- Apple Feedback addresses
 - opentpt@seeding.apple.com
 - rhapsody-dev-feedback@apple.com



Your Evangelist !!!

Thomas Weyer

N&C Evangelism

Apple Computer, Inc.

3 Infinite Loop, MS 303-2E

Cupertino, CA 95014

Phone: (408) 974-5017

Fax: (408) 974-1211

Email: weyer@apple.com

<http://www.devworld.apple.com>

[dev/opentransport/](http://www.devworld.apple.com/dev/opentransport/)





Single Link Multihoming

David Kim

Senior Software Engineer
MacOS Networking

Introduction

- Ability to configure additional IP Secondary addresses with IP Primary address active
- Meet needs exhibited by ISPs primary for the purpose of Vanity IP addresses multiple hosting by Internet servers
- Not for End User



Configuration

- IP Primary address is configured using the TCP/IP control panel
- IP Secondary addresses are configured using entries in a text file named “**IP Secondary Addresses**”, which resides in the System Preferences folder
- Addresses are not validated, duplicate addresses will be reported
- The number of addresses is only limited by available memory



Support

- New API to retrieve IP Secondary addresses
 - pascal OSStatus
OTInetGetSecondaryAddresses
(InetHost* addr, UInt32*
count, SInt32 val);
- Primary address must be configured
Manually



Support

- IP Secondary addresses:
 - Do not have to be configured in a contiguous block
 - Will use the same subnet as the Primary address



IP Secondary Addresses Format

- One IP address per line
- Example:

```
;list of IP Secondary addresses  
15.0.0.2  
15.1.1.5  
15.3.4.5  
;15.1.2.0
```

- Use of “;” on a line indicates the line is ignored, as in Hosts file



Developer Tips

- New `fIPSecondaryCount` field is added to `InetInterfaceInfo` structure
- Call `OTInetGetInterfaceInfo` for the `fIPSecondaryCount`
- Call `OTInetGetSecondaryAddresses` for all the IP Secondary addresses
- Call `OTBind` with a specific IP address to listen on that address





Demo



OT API Change

John Comiskey

Senior Software Engineer
MacOS Networking

Overview

- These APIs are new in OT 1.1.1.
- This is a subset, refer to the Development Notes for more complete information.



New DLPI Template

- `dlpiether.c` is a new DLPI driver template from Mentat that replaces previous template from Apple
- Supports version 2 of the DLPI specification
- Included in SDK along with detailed documentation



OTSetMemoryLimits and OTSetServerMode

- Force OT memory pool to grow immediately and never shrink
- Useful for servers that don't want to wait for the pool to grow, which can only happen at System Task time
- Prototypes for `OTSetMemoryLimits` and `OTSetServerMode` are not in any header files to prevent them from being abused



OTSetMemoryLimits

- OSStatus

```
OTSetMemoryLimits(size_t  
growSize, size_t maxSize);
```

- Pool will immediately grow by growSize bytes and will never shrink
- The maximum pool size is set to maxSize
- Default maxSize for OT is 10% of physical memory



OTSetServerMode

- `void OTSetServerMode();`
 - Pool will immediately grow by 2MB and will never shrink
 - Does not change the maximum pool



OTUseSyncIdleEvents

- Used to enable or disable `kOTSyncIdleEvent` for your endpoint
- When enabled your notifier will be called with a `kOTSyncIdleEvent` repeatedly while OT is waiting for a sync call to complete
- It is safe to call `YieldToAnyThread` at this point since it is only called at S Task time. This is the only notifier that can safely be used to call `YieldToAnyThread`



OTUseSyncIdleEvents

- The notifier for the `ProviderRef` will receive `kOTSyncIdleEvent` while the `ProviderRef` is blocked by OT

- OSStatus

```
OTUseSyncIdleEvents(ProviderRef  
ref, Boolean useEvents);
```

- `ref` is the `ProviderRef` (i.e., endpoint that wants to receive `kOTSyncIdleEvent`)
- `useEvents` is true to enable `kOTSyncIdleEvent`, false to disable it



OTEnterNotifier and OTLeaveNotifier

- Used to prevent your notifier from being entered
- Useful for critical sections outside of the notifier



OTEnterNotifier and OTLeaveNotifier

- Boolean

```
OTEnterNotifier(ProviderRef  
ref);
```

- returns true if successful

- returns false if the notifier is currently entered

- void

```
OTLeaveNotifier(ProviderRef  
ref);
```



OTGate API

- Used to serialize the the processing data or events
- Common use is to serialize OT notif events across multiple endpoints to critical section issues
- Consists of `OTInitGate`, `OTEnterGate` and `OTLeaveGate`
- See “Open Tpt Module Dev. Note”



OTGate API

- Client example:
 - Serialize OT events with API calls that your code handles from above
 - HTTP implemented as an OT client would need to serialize OT events with calls to its API
- Module example:
 - When the sync queue level is less restrictive than `SQLVL_MODULE`





Open Transport Performance

Quinn “The Eskimo!”
Developer Technical Support

Overview

- Why you *don't* need this talk!
- General rules for optimization
- Measuring your performance
- OT-specific optimization tips
- General Mac OS optimization tips



Why You Don't Need This T

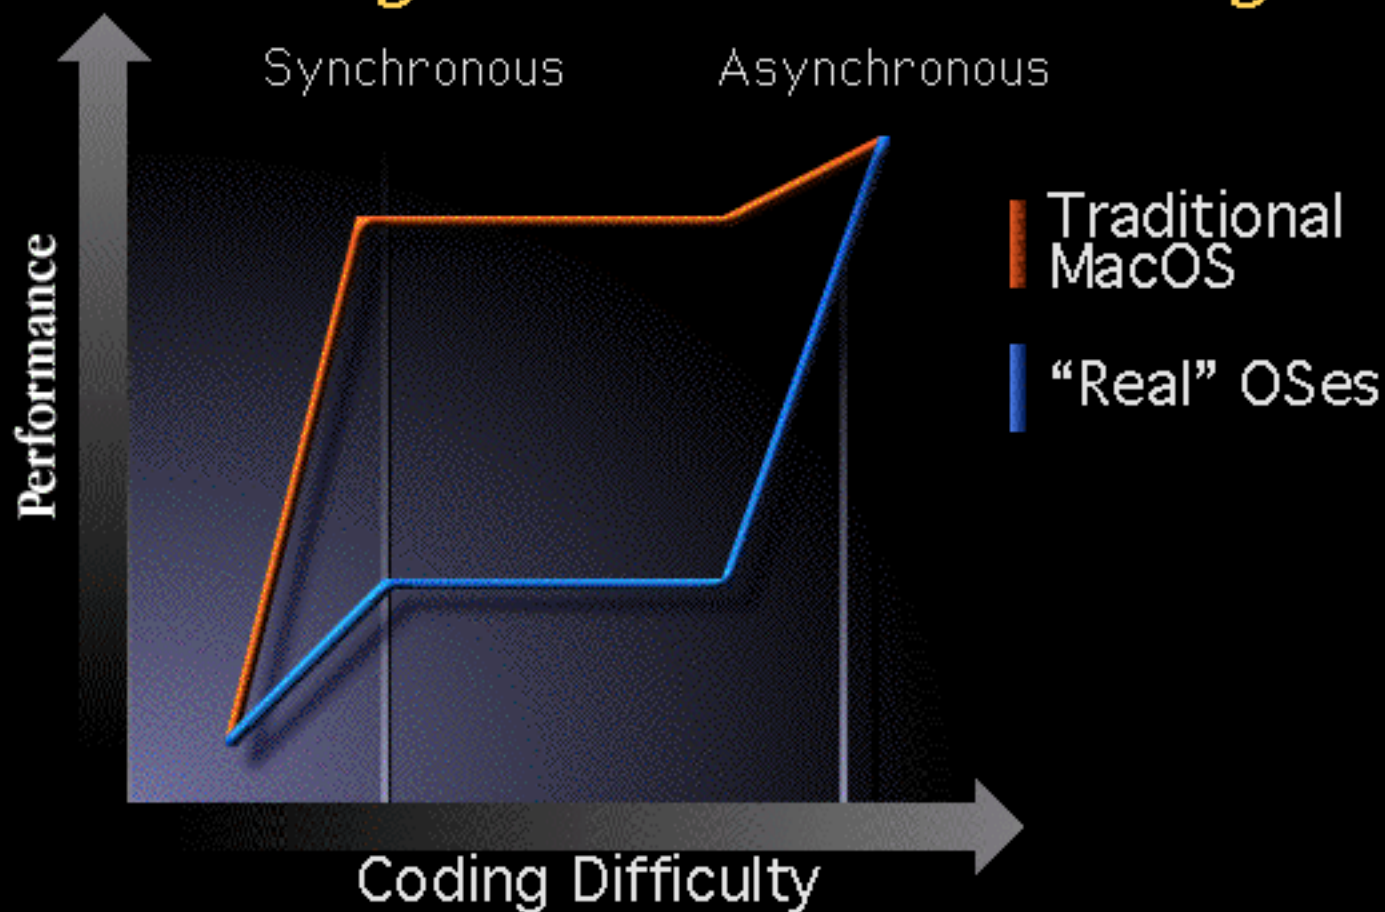
Avoid Premature Optimization

- As a DTS engineer, I see this problem the time
- Bo3b's Rule: I'd rather be windsurfing
- It's hard to get 100% performance traditional Mac OS
 - So don't try unless you are 100% certain that your product needs it to succeed
 - e.g., NewsWatcher, a very successful Internet application, uses polled I/O



Quinn's Bogus Chart

Networking Performance vs Coding Effort



General Optimization Tenets

- Make it work, then make it go fast
 - Reliability is more important than speed
- Algorithms have the most effect
 - If you control both ends of the connection, design your protocol for speed
 - Unfortunately, this is often not the case
- Decide on good performance metrics
- Set a performance goal before you start
- Stop when you hit it
- Profile before you change code



Algorithmic Examples

Change Your Approach, Go Much Faster

- AppleShare IP
 - You could tweak the small things until the end of time, but you would never get the speed benefits that switching to a faster underlying protocol yields
- FTP client
 - Using polled I/O for the command connection, notifiers for the data connections



Performance Metrics

- Highly application specific
- Maybe your program has an industry standard
 - e.g., WebStones?
- Otherwise, you have to develop your own metrics
- Don't forget subjective metrics



Profile Before You Start

- Otherwise you risk wasting time optimizing the wrong thing
 - e.g., server might be file system bound or network bound—if it's file system bound you're in the wrong talk!
- Instrumentation SDK
 - Allows you to add trace points to your code
 - Unlike normal profilers, these trace points aren't based on routine entry/exit
 - Works well for asynch I/O environments



Optimizing Your Network Co

- The next few slides contain specific suggestions
- Each suggestion has a corresponding caution
- You must decide which optimizations are suitable for your application



OT Optimization #1

Program Structure

- Do everything in your notifier...Why
 - `WaitNextEvent` takes forever to return
 - 10 ticks *is* forever on a 100Mbps network
 - Notifiers offer important re-entrancy guarantees
 - OT offers interrupt-safe memory management
- However, there are alternative program structures that may be fast enough



Caution #1

Program Structure

- The listen/accept hand-off sequence for OT is particularly scary
 - See Technote 1059 for details
 - The `tilisten` module helps
- You must also take care to avoid `kOTStateChangedErr`
 - Use `OTEnterNotifier` to sync foreground and notifier threads
 - Wait for completion events before starting new operations



OT Optimization #2

Cache and Reuse Endpoints

- It takes longer to create/destroy endpoints than to bind/unbind them
- Thus, if your program uses a lot of endpoints (e.g., a web server), you gain performance by keeping a cache of unbound worker endpoints
- Don't bind worker endpoints, `OTAConnect` will do that for you



Caution #2

Cache and Reuse Endpoints

- This technique has very few drawbacks
- You may find situations where the `OTUnbind` fails—in that case just close and reopen the endpoint
- Also remember that each endpoint consumes memory that might be used elsewhere



OT Optimization #3

Avoid Copying Data

- OT is very efficient about data copy
 - In the best case, OT can receive and transmit without any copying
 - But you must explicitly enable this a
- No-copy `OTRcv`
 - Returns pointer to OT buffers, you m
 - return it with `OTReleaseBuffer`
- Ack Sends mode
 - You must keep data around until `T_MEMORYRELEASED`



Caution #3

Avoid Copying Data

- For no-copy receives, make sure you return the receive buffer to OT as o as possible
 - In many cases, this buffer will be the driver's DMA buffer
- Ergo, don't modify the buffer!



Caution #3 (cont.)

Avoid Copying Data

- Ack Sends will complicate your code because you must track memory allocation yourself
 - e.g., calling `CloseOpenTransport` with outstanding Ack Sends is fatal
- Under certain circumstances, OT may modify your data buffer



OT Optimization #4

Use OT Utilities Where You Can

- OT LIFO lists are better than the traditional OSUtils *Enqueue/Dequeue* routines
- OT's memory management is optimized for networking
- *OTTimeStamp* is the best timestamping on the traditional Mac
- OT's deferred tasks will adjust to system changes better than *DTInst*



Caution #4

Use OT Utilities Where You Can

- No caution needed. Just do it!



OT Optimization #5

Server Mode

- `OTSetMemoryLimits` allows you to increase the amount of memory OT use for buffers
- Useful for servers that are fielding many simultaneous connections
- If you just want to push the maximum number of bits through a *single pipe*, most probably won't help



Caution #5

Server Mode

- `SetMemoryLimits` is a case where a performance metric is very useful
 - Giving more memory to OT may improve your performance
 - ...or it may do nothing
 - You need to measure it for your particular application
- Definitely not recommended for client software



OT Optimization #6

Protocol Tweaks

- **XTI_RCVBUF**
 - Alters the endpoint's receive buffer size
 - Affects the TCP “window size”
- **XTI_SNDBUF**
 - Alters the endpoint's send buffer size
- TCP protocol has a lot of (currently private) knobs you can twiddle



Caution #6

Protocol Tweaks

- Like `SetMemoryLimits`, these techniques require intelligent use
- Changing these values can improve, worsen, or have no effect
- The effect may even vary from site to site
- My recommendations:
 - Test in production environments to find sensible defaults
 - If you find great variance, put it under user control



General Mac OS Optimizatio

- Try to avoid the file system
 - Cache everything you can
 - When you can't, call the file system asynchronously
 - And call it for big chunks of data
 - Double buffer file and network I/O
 - See DTS Technote FL 16 “File Management Performance and Caching”
- Avoid taking MixedMode hits
 - Don't disable interrupts



Virtual Server Sample

- HTTP-like sample written by the OT TCP/IP lead
- Approximately 400 connections per second on 200 MHz 604
- Very useful sample if you need speed
- Shows good techniques for most of optimizations in this talk
- Also shows how to workaroud a number of OT 1.1.x bugs



Wrap Up

- Think before you optimize
- Define your performance goals before you begin
- OT can be very fast, if you're prepared to do the work
- OT Virtual Server is the place to start



Suggested Reading

- OT Web Page
 - <http://devworld.apple.com/dev/opentransport/>
- Inside Macintosh: Open Transport
- Technotes
 - Technote 1059 “On Improving Open Transport Network Server Performance”
 - Technote FL 16 “File Manager Performance and Caching”
- Instrumentation SDK—ETO 23





Q&A

The background features a dark, textured surface with a glowing blue and purple sphere in the center. A white Apple logo is positioned at the top of the sphere. The text "Worldwide Developers Conference" is overlaid on the image. "Worldwide" and "Conference" are in a gold, serif font, while "Developers" is in a white, serif font enclosed in a white rectangular box. The overall aesthetic is futuristic and tech-oriented.

Worldwide

Developers

Conference