

Obtaining Gated

- Anonymous FTP

gated.cornell.edu

pub/gated/gated-2.0.1.14.tar.Z

pub/gated/gated-alpha.tar.Z

pub/lists/gated-people

- Electronic mail

gated-people-archive@gated.cornell.edu

- Subject:

- * send gated/gated-2.0.1.14.tar.Zus*

- * send gated/gated-alpha.tar.Zus*

- * help

- * send Contents

Work in Progress ...

- Integrated stack IS-IS
 - University of Wisconsin WISIS
- Dynamic Interfaces
 - *Pop-up* Interfaces
 - * SunOS, Ultrix SLIP
 - Unavailable Interfaces
- IDRP
 - Dave Katz (Merit)
- Router Discovery
 - Erik Nordmark (SunSoft)

Work in Progress

- BSD 4.3 Reno radix tree
 - Variable length masks
 - Variable length addresses (ISO)
 - Ordered retrieve (SNMP)

Parser

- Interface reference with wildcard

le → *le0*, *le1*, *le2*

- Trace options per task

Tasks ...

- Access to sockets

- Common socket options and ioctls

task_set_option()

- * Multicast parameters

TASKOPTION_MULTI_ADD

TASKOPTION_MULTI_DROP

TASKOPTION_MULTI_IF

TASKOPTION_MULTI_LOOP

TASKOPTION_MULTI_TTL

- * Receive dest address (Reno)

TASKOPTION_RCVDESTADDR

- * Use loopback (routing socket)

TASKOPTION_USELOOPBACK

Tasks

- Memory management
 - System memory management

task_mem_malloc()

task_mem_calloc()

task_mem_realloc()

task_mem_free()

- Optimized for control blocks

task_block_init()

task_block_alloc()

task_block_free()

SNMP

- Via SMUX protocol
 - ISODE 7.0 w/SNMP UPGRADE 1
 - Freely available
 - Robust
 - Compiled in MIB
 - Traps
 - Writable variables

Routing Table ...

■ Preferences

0	<i>interface</i>
10	<i>OSPF (intra/inter area)</i>
20	<i>default (EGP/BGP)</i>
30	<i>redirects</i>
40	<i>'route' (routing socket)</i>
50	<i>SNMP</i>
60	<i>static</i>
90	<i>HELLO</i>
100	<i>RIP</i>
110	<i>OSPF (ASE)</i>
150	<i>BGP</i>
200	<i>EGP</i>

Routing Table ...

■ Route Locking

- Protocols flag **routes** they announce
- **Routes** not released until flag reset
- Holddowns implemented by protocols that require them (RIP, HELLO)
- Non-holddown protocols not restricted by holddowns (BGP, EGP, OSPF)

■ Policy independence

- All routes are retained (except RIP, HELLO and redirects)
- Reconfiguration results in minimal impact

Routing Table

- Event driven
 - Policy evaluated by protocols when changes occur
 - Steady state processing minimal (just packet stuffing)
- Kernel routing table
 - Implemented as a protocol
 - Always called first

HELLO

- Proper subsuming of host routes
- **Split horizon w/o poison reverse**
 - Reduces routing overhead

RIP

- RIP version 1 (RFC-1058)
- Proper subsuming of host routes
- **Split horizon w/o poison reverse**
 - Reduces routing overhead
- Graceful shutdown
 - Avoids **black holes**

OSPF

- Version 2 (RFC-1247)
- Multicast support (some systems)
- MIB support
- Not reconfigurable

BGP

- BGP version 2 (RFC-1163)
- BGP version 3 (RFC-????)
- AS path pattern matching (RFC-1164)
- Dennis Ferguson (CA*NET)
 - Entering production

Version 3.0

- BGP version 2/3 (RFC-1163/RFC-????)
- EGP version 2 (RFC-904)
- OSPF version 2 (RFC-1247)
- RIP version 1 (RFC-1058)
- HELLO (Fuzzballs)
- SNMP via ISODE and SMUX
- Routing socket (some systems)
- Variable subnet masks (some systems)
- Integrated stack IS-IS
- IDRPF (ISO Inter-Domain Routing Protocol)

Routing Table ...

- Route selection

- Degree of Preference

- * **0** most - **255** least preferred

- 0** ***Interface***

- 10** ***Default***

- 20** ***Redirect***

- 50** ***Static***

- 90** ***HELLO***

- 100** ***RIP***

- 150** ***BGP***

- 200** ***EGP***

- Metric

- * **Same** protocol and AS

- Gateway with lowest IP address

Routing Table

- Multiple **routes** (RIP, EGP, ...) per **destination** (network/mask)
- Static Interface routes
 - Multiple logical networks
 - Only allows access to gateways

Tasks ...

- Reconfiguration

- Preconfig cleanup

task_cleanup

- After config setup

task_reinit

- After config policy re-evaluation

task_newpolicy

- Limited subprocess support

task_fork()

task_kill()

- Used by status dump

/var/tmp/gated_dump

Tasks ...

- Access to sockets ...
 - Common socket options and ioctls
 - task_set_option()*
 - * Receive/Send buffer size
 - TASKOPTION_RECVBUF***
 - TASKOPTION_SENDBUF***
 - * Broadcast
 - TASKOPTION_BROADCAST***
 - * Bypass routing table
 - TASKOPTION_DONTROUTE***
 - * Linger after close
 - TASKOPTION_LINGER***

Tasks ...

- Access to sockets ...
 - Ready for write
task_write
task_send_packet()
 - Exception
task_except
 - Connection complete
task_connect
 - Incoming connection
task_accept

Tasks ...

- Access to sockets

task_get_socket()

task_set_socket()

task_reset_socket()

task_close()

- Establishing connections

task_addr_local()

task_addr_remote()

- Packet ready/Data ready

task_recv

task_receive_packet()

Tasks

- One or more per protocol

task_alloc()

task_create()

task_delete()

- Routing table changes

task_flash

- Interface changes

task_ifchange

- Up/Down
- Add/Delete

Timers

- One second resolution
- One-shot or repeating
- Corrected for drift (**anti-boxcar**)
- Separate queues for
 - Inactive
 - Active (in order)

timer_create()

timer_delete()

timer_set()

timer_interval()

timer_reset()

Family common

- Checksum

inet_cksum()

- Masking

inet_netmask()

inet_wholenetof()

- Classes (Internet)

inet_class()

Parser

- Expandable Lex/Yacc (Flex/Bison)
- Improved diagnostics
- Facilities for common types
 - IP addresses/masks
 - Time (ss, mm:ss, hh:mm:ss)
 - metrics (w/limit checking)
 - interface references
 - gateways (host on local network)
- Structured grammar (C-like syntax)
- Nested configuration files
 - *%include*

Tracing

- Uniform trace file and syslog support
 - %m format for system errors
 - %A format for addresses (IP, ISO)
 - %T format for time

tracef()
trace()

- Many levels

internal, external, route, kernel
parse, lex, config
task, timer, mark, nostamp
update, protocol

Facilities

- Tracing
- Parser
- Family common
- Timers
- Tasks
- Routing table

Redirects

- Configuration option to disable
 - Deletes and reinstalls old route
 - Default when supplying RIP/HELLO

RIP

- Version 1 (RFC-1058)
 - Does not reject packets based on zero fields
- Does flash updates correctly

HELLO

- Supports subnetwork and hosts
- Does incremental updates correctly

EGP

- Version 2 (RFC-904)
- **Complete** implementation
- Allows grouping of peers for maxup
- Maximum packet size configurable
- HELLO and POLL timers configurable

BGP

- Version 1 (RFC-1105)
- Route selection on AS path (RFC-1106)
not yet implemented

Version 2.1

- BGP version 1 (RFC-1105)
- EGP version 2 (RFC-904)
- RIP version 1 (RFC-1058)
- HELLO (Fuzzballs)
- Configuration parser
- Multi-way policy

Design Goals

- Modularity
 - Protocols are independent
 - Selective compilation
- Capability
 - More configurable than routers
- Simplicity
 - Minimal configuration for simple case
 - Easy to configure advanced features

What is Gated?

- Routing daemon for **BSD** based systems
- Initially as **glue** for the 56kB NSFnet
- Originally based on BSD **routed** and Kirton's **egpup**
- Partially funded by NSF
- Model implementation of routing protocols
- Routing protocol development platform

GateDaemon (Gated)

Information Technologies/
Network Resources
Cornell University
143 Caldwell Hall
Ithaca, NY 14853-2602

gated-people@gated.cornell.edu
gated-people-request@gated.cornell.edu