

## Chapter 8

### gdc dump

#### 8.1 Purpose of a GateD Dump File

The purpose of a GateD dump file is to create an ASCII representation of GateD's internal state. A GateD dump file is not a core file and cannot be used with a debugger.

Most protocols and pseudo-protocols in GateD implement a "dump method" that is called to dump data specific to the module. In addition, the entire routing table is dumped (including an ASCII representation of the radix trie(s)) along with other pieces of core state. This data is a "snapshot" of the internal state at some point in time. Due to resource scheduling, this may not be the exact moment at which the dump signal was received.

#### 8.2 When to Create a Dump File

Create a dump if GateD is running with a known or suspected error.

The GII and SNMP tools provide alternatives to a GateD dump. GII provides protocol info, a view of the current state of GateD, and a routing table listing.

#### 8.3 How to Create a Dump

Create a GateD dump file using one of the following:

- `gdc dump` command
- `SIGINT` signal from UNIX system

To create a dump using the dump command, type

```
gdc dump
```

To create a dump with `SIGINT`, send a `SIGINT` signal to GateD. GateD forks a child process to perform the dump while the parent continues on as before. The `kill(1)` program is typically used to send signals to processes.

#### 8.4 Finding a Dump

The default file location for a GateD dump is `/var/tmp/gated.dump`. This path is determined at compile time; the alternative location is `/usr/tmp/gated_dump`.

#### 8.5 Format of GateD Dump File

GateD dumps include information about the following:

- general GateD state
- tasks
- memory usage and statistics

- interfaces
- martians
- routing table(s) (including info about configured statics)
- policy
- protocol-specific data

### 8.5.1 General Information in Dump

The first section of the dump will contain general information such as:

- GateD version
- date of dump
- operating system
- tracing flags set

In the following example,

- the GateD version is 8.0.0
- the date the dump was generated is January 18, 2000
- the operating system used is *SunOS 5.6*
- the tracing flags set are **a11**
  - the file is open
  - 12 tasks reference the file
  - the current logfile size is 88,542 bytes
  - the maximum size is 102,4000 bytes
  - 2 old versions of the file are kept

---

```
gated[12843] version 8-0-0 memory dump on paris.merit.edu at Tue Jan 18
01:46:41 2000
```

```
SunOS 5.6 Generic_105181-17 paris.merit.edu sun4u sparc
```

```
Started at Tue Jan 18 01:42:30 2000
```

```
Tracing:
```

```
Global: all
```

```
Files:
```

```
    /var/tmp/gated.log:
```

```
opened refcount 12 size 88542 limit 1024000 files 3
```

---

### 8.5.2 Task State

The task section of the dump contains information about the task state, such as:

- send/rcv buffer size
- task name

- task priority
- tracing options for task
- trace file for task
- timers associated with task

In the following example:

- the send buffer size is 16384, and the receive buffer size is 8,192 bytes
- the task names are IF, INET, Aggregate, RT, ICMP, and BGP.0.0.0.0+179
- the task priorities for each task are, respectively, 10, 15, 20, 20, 30, and 40
- the tracing options are: general, state, policy, task, and timer for all names except BGP, which has only packets
- the trace files are all : /var/tmp/gated.log
- the only timer is associated with IF and is IF\_AGE
  - the time of last firing is 01:45:30.259 am; next firing is at 01:48:30.259 AM

---

Task State: <>

Send buffer size 16384 at 150000

Recv buffer size 8192 at 14A000

Tasks (12) and Timers:

IF                    Priority 10        RtProto Direct  
 Trace options: general state policy task timer  
 Trace file: /var/tmp/gated.log size 1024000 files 3

IF\_AGE <oneshot>

last: 01:45:30.259        next: 01:48:30.259

INET                Priority 15        Socket 4        RtProto INET  
 Trace options: general state policy task timer  
 Trace file: /var/tmp/gated.log size 1024000 files 3

Aggregate          Priority 20        RtBit 1  
 Trace options: general state policy task timer  
 Trace file: /var/tmp/gated.log size 1024000 files 3

RT                   Priority 20  
 Trace options: general state policy task timer  
 Trace file: /var/tmp/gated.log size 1024000 files 3

ICMP                Priority 30        Proto 1        Socket 6  
 Trace options: general state policy task timer

```
BGP.0.0.0.0+179 Priority 40      Port    179      Socket  8      RtProto BGP
Trace options: packets
Trace file: /var/tmp/gated.log  size 1024000 files 3
```

---

### 8.5.3 Socket Read Routines

The next section of the dump will contain socket read information such as:

- read routines (in the example: KRT, ICMP, GII\_LISTEN)
  - low priority read routines (BGP.0.0.0.0+179)
  - file descriptors (files and sockets)
- 

Socket read routines:

```
1      KRT
6      ICMP
7      GII_LISTEN.0.0.0.0+616 (accept)
```

Socket low priority read routines:

```
8      BGP.0.0.0.0+179 (accept)
```

File Descriptors (max 8):

```
0
1      Task: KRT      read rqueue
2      Task: KRT      File: krt_ifread_task
3      File: /etc/gated.pid
4      Task: INET
5
6      Task: ICMP      read rqueue
7      Task: GII_LISTEN.0.0.0.0+616  accept rqueue
8      Task: BGP.0.0.0.0+179          accept
```

---

## 8.5.4 Memory Usage in Dump

This section of the dump will contain memory usage information for each task block such as:

- allocation size (8,192 bytes in the example)
- number of free blocks (935 blocks)
- the last block allocated (001483B0)
- data types for which memory is allocated (runt, krt\_remnant\_rt, vtxlist\_t, nospf\_route\_link\_t)

---

Task Blocks:

Allocation size: 8192

Size: 8 N\_free: 935 LastAlloc: 001483B0

runt	Init:	5	Alloc:	0	Free:	0	InUse:	0
krt_remnant_rt	Init:	1	Alloc:	5	Free:	5	InUse:	0
vtxlist_t	Init:	1	Alloc:	0	Free:	0	InUse:	0
nospf_route_link_t	Init:	1	Alloc:	0	Free:	0	InUse:	0
iflist_t	Init:	1	Alloc:	0	Free:	0	InUse:	0
sockaddr_un.in	Init:	1	Alloc:	309	Free:	217	InUse:	92
asmatch_t	Init:	1	Alloc:	0	Free:	0	InUse:	0

---

## 8.5.5 Interface - Task and Logical Addresses

The interface section of the dump will contain interface information for each task such as:

- physical interfaces assigned and up (in the example, IF has 2 physical interfaces assigned and 2 up)
- INET protocol addresses assigned and up (2 assigned and 2 up)
- local and physical addresses to which the interface connects (local: 120.0.0.1, 198.108.60, 192.168.10, and 0.0.0.142; physical: 8:0:20:a8:8c:60 )
- interface names (lo, lo0, hme, hme0, hme1)
  - index number (for lo:1)
  - change (for lo:<> )
  - state (for lo: <up loopback multicast>)
  - refcount (for lo:2)
  - up-down transitions (for lo:0 120.0.1)
  - metric (for 120.0.1:0)
  - MTU (for 120.0.1:1472 )
  - refcount (for 120.0.1:3)
  - preference when interface is up (for 120.0.1:0)
  - preference when interface is down (for 120.0.1:120)

- change (for 120.0.1: <> )
- state (for 120.0.1: <up loopback multicast>)
- subnet mask (for 120.0.1: 255.255.255)

---

Task IF:

Physical interfaces: 2           Up: 2  
INET protocol addresses: 2       Up: 2

Addresses:

120.0.1  
P2P 0   Loop 1   Total 1   Refcount 2   Route: not installed  
198.108.60.142  
P2P 0   Loop 0   Total 1   Refcount 1   Route: not installed  
192.168.10.142  
P2P 0   Loop 0   Total 1   Refcount 1   Route: not installed

Local addresses:

127.0.0.1  
P2P 0   Loop 1   Total 1   Refcount 1   Route: not installed  
198.108.60.142  
P2P 0   Loop 0   Total 1   Refcount 1   Route: not installed  
192.168.10.142  
P2P 0   Loop 0   Total 1   Refcount 1   Route: not installed  
0.0.0.142  
P2P 0   Loop 0   Total 0   Refcount 2   Route: not installed

Physical addresses:

802.2 8:0:20:a8:8c:60

Refcount 2

Names:

lo  
                  Refcount 1

lo0  
                  Refcount 1

hme  
                  Refcount 2

hme0  
                  Refcount 1

---

```
hme1
    Refcount 1
Interfaces:
  lo0      Index 1      Change: <>      State:<up loopback multicast>
    Refcount: 2      Up-down transitions: 0

    120.0.1
      Metric: 0      MTU: 1472
      Refcount: 3      Preference: 0      Down: 120
      Change: <>      State: <up loopback multicast>
      Subnet Mask: 255.255.255.255

hme0      Index 2      Address 802.2 8:0:20:a8:8c:60      Change: <>      State:
    Refcount: 2      Up-down transitions: 0

    198.108.60.142
      Metric: 0      MTU: 1436
      Refcount: 3      Preference: 0      Down: 120
      Change: <>      State: <up broadcast multicast>
      Broadcast Address: 198.108.60.255
      Subnet Number: 198.108.60      Subnet Mask: 255.255.255

hme1      Index 3      Address 802.2 8:0:20:a8:8c:60      Change: <>      State
:<up broadcast multicast>
    Refcount: 2      Up-down transitions: 0

    192.168.10.142
      Metric: 0      MTU: 1436
      Refcount: 16      Preference: 0      Down: 120
      Change: <>      State: <up broadcast multicast>
      Broadcast Address: 192.168.10.255
      Subnet Number: 192.168.10      Subnet Mask: 255.255.255
```

---

## 8.5.6 INET - Martians

The martians section of the dump will contain INET task information, which is where the martians configuration is stored. Martians are configured using the 'martians' clause. (See "Martian Syntax" on page 30 in *Configuring GateD* for more information.) The dumped data includes the state of:

- IP forwarding (in the example, INET has 0 IP forwarding)
- UDP checksums (1)
- reject address (127.0.0.1)
- blackhole address (127.0.0.1)
- autonomous system (201)
- router ID (0.0.0.142)
- martians (0.0.0.0, 127, 255.255.240)

---

Task INET:

IP forwarding: 0                      UDP checksums: 1

Reject address: 127.0.0.1      Blackhole address: 127.0.0.1

Autonomous system: 201

Router ID: 0.0.0.142

Martians:

0.0.0.0	mask 0.0.0.0	Exact
0.0.0.0	mask 255	
127	mask 255	restrict
255.255.240	mask 255.255.240	restrict

---

## 8.5.7 Route Table (RT) Information

This section of the dump contains information about configured static routes as well as the state of the RIBs (Routing Information Base(s)) for the various address families. Only 'inet' will be shown in the following examples. The routes are printed in an ASCII representation of the radix trie as well as in a detailed listing on a per-destination basis.

---

Task RT:

Static routes for family INET: (\* indicates gateway(s) in use)

Radix trie for inet (2) inodes 27 routes 15:

```
          +-24+--{223.145.154
        +-22+
        | | +-24+--{223.145.153
```

```

| +-23+
| +-24+---{223.145.152
+-20+
| | +-24+---{223.145.151
| | +-23+
| | | +-24+---{223.145.150
| | +-22+
| | | | +-24+---{223.145.149
| | | +-23+
| | | +-24+---{223.145.148
| +-21+
| | +-24+---{223.145.147
| | +-23+
| | | +-24+---{223.145.146
| +-22+
| +-24+---{223.145.145
+--8+
| | +-24+---{223.3
| | +-15+
| | | +-24+---{223.2
| +-14+
| +-24+---{223.1
+--0+---{0.0.0.0
| +--8+---{127

```

It also includes static route detail information, such as:

- net mask set for the route (in the example, mask 0.0.0.0 is set for route 0.0.0.0; mask 255 is set for route 127)
- preference set for the route (preference 60 for route 0.0.0.0; 0 for route 127)
- which gateway/interface is used for the route (gateway 198.108.60.1 is set for route 0.0.0.0; interfaces 127.0.0.1 is set for route 127)

```
0.0.0.0 mask 0.0.0.0 preference 60 state <int retain gateway> Gate-
way 198.108.60.1(*)
```

```
127 mask 255 preference 0 state <noadvise int retain reject>
Interfaces 127.0.0.1
```

```
223.1      mask 255.255.255 preference 60  state  Gateway 192.168.10.1(*)
```

```
223.2      mask 255.255.255 preference 60  state <int gateway> Gateway
192.168.10.1(*)
```

```
223.3      mask 255.255.255 preference 60  state <int gateway> Gateway
192.168.10.1(*)
```

---

The next portion of the dump shows static gateway information, such as which gateways and tasks are referenced by static routes. (In this example, no routes reference either of the gateways. Task RT is referenced by both of the routes.)

---

Gateways referenced by static routes:

```
198.108.60.1
           routes: 0 task: RT
192.168.10.1
           routes: 0 task: RT
```

---

This is followed by static route announce bit information, such as bit allocations, as shown in the example below.

---

Bit allocations:

```
1      Aggregate
2      KRT      byte index: 0  length: 4
```

---

Static route mask and address information are shown in the example below, such as:

- family
- address
- length
- mask

---

Masks and addresses:

Family	Address	Length	Mask
inet	1317B0	0	0.0.0.0
inet	1317B8	1	128
inet	1317C0	2	192
inet	1317C8	3	224
inet	1317D0	4	240
inet	1317D8	5	248

```

inet    1317E0  6      252
inet    1317E8  7      254

```

---

The sixth section of the dump will also include a graph of the real (nonstatic) routes as shown in the example below.

---

Routing Tables:

Radix trie for inet (2) inodes 32 routes 18:

```

          +-24+---{223.145.154
        +-22+
        | | +-24+---{223.145.153
        | +-23+
        | +-24+---{223.145.152
    +-20+
    | | +-24+---{223.145.151
    | | +-23+
    | | | +-24+---{223.145.150
    | | +-22+
    | | | | +-24+---{223.145.149
    | | | +-23+
    | | | +-24+---{223.145.148
    | +-21+
    | | +-24+---{223.145.147
    | | +-23+
    | | | +-24+---{223.145.146
    | +-22+
    | +-24+---{223.145.145
    +-8+
    | | +-24+---{223.3
    | | +-15+
    | | | +-24+---{223.2
    | +-14+
    | +-24+---{223.1
    +-3+
    | | +-24+---{198.108.60
    | +-5+

```

```
      |      +-24+--{192.168.10
+--0+--{0.0.0.0
      |      +--8+--{127
      |      +-32+--{127.0.0.1
```

+ = Active Route, - = Last Active, \* = Both

---

The next section of the dump shows real route detail information, such as:

- for all the routes:
    - number of destinations (in the example that follows, 18 destinations exist for these routes)
    - number of routes (18)
    - number of holddowns (0)
    - number of deleted routes (10)
    - number of hidden routes (0)
  - for each route:
    - net mask (for route 0.0.0.0, the mask is 0.0.0.0)
    - number of entries (for route 0.0.0.0, the number of entries is 1)
    - announce (1)
    - TSI (none)
    - instability histories (none)
    - \*static preference (60)  
**Note:** The asterisk (\*) indicates that route is active.
    - next hop (198.108.60.1)
    - interface (198.108.60.142(hme0))
    - state bits set<int active retain gateway>
    - age (4:11)
    - metric (0)
    - metric2 (0)
    - tag (0)
    - task (RT)
    - announcement bits(1) (2-KRT)
    - AS path (IGP (Id 1))
-

---

Routing table for inet (2):

Destinations: 18            Routes: 18  
Holddown: 0        Delete: 10        Hidden: 0

0.0.0.0        mask 0.0.0.0

entries 1        announce 1

TSI:

Instability Histories:

\*Static Preference: 60

NextHop: 198.108.60.1    Interface: 198.108.60.142(hme0)

State: <int active retain gateway>

Age: 4:11        Metric: 0        Metric2: 0        Tag: 0

Task: RT

Announcement bits(1): 2-KRT

AS Path: IGP (Id 1)

127            mask 255

entries 1        announce 1

TSI:

Instability Histories:

\*Static Preference: 0

NextHop: 127.0.0.1    Interface: 127.0.0.1(lo0)

State: <noadvise int active retain reject>

Age: 4:11        Metric: 0        Metric2: 0        Tag: 0

Task: RT

Announcement bits(1): 2-KRT

AS Path: IGP (Id 1)

127.0.0.1    mask 255.255.255.255

---

## 8.5.8 Example Protocol-Specific Data: BGP Peers and Groups

The protocol-specific section of the dump includes information about all protocols that implement a dump method. The example shown here shows BGP peer and group information, such as:

- task name
  - task groups
  - task total peers
  - task active incoming peers
  - task free peers
  - task free groups
  - task state
  - task autonomous system (AS) number
  - import controls for each AS
  - net mask for each AS

---

```
Task BGP.0.0.0.0+179:
Groups: 1   Peers: 1 (1 configured)   Active Incoming: 0
Free Peers: 0Free Groups: 0           State: Listening
  AS 202:
    Import controls:
      0.0.0.0           mask 0.0.0.0

  AS 201:
    Import controls:
      0.0.0.0           mask 0.0.0.0

  AS 202:
    Export controls:
      Protocol BGP  as 202
        0.0.0.0           mask 0.0.0.0
      Protocol BGP  as 201
        0.0.0.0           mask 0.0.0.0
      Protocol Static
        0.0.0.0           mask 0.0.0.0

  AS 201:
    Export controls:
      Protocol BGP  as 202
```

```

          0.0.0.0          mask 0.0.0.0
    Protocol BGP  as 201
          0.0.0.0          mask 0.0.0.0
    Protocol Static
          0.0.0.0          mask 0.0.0.0

```

Task BGP\_202.192.168.10.101+179:

```

Peer: 192.168.10.101+179  Local: 192.168.10.142  Type: External
State: Active  Flags: <>
Last State: Idle          Last Event: Start          Last Error: None
Options: <>

```

Task BGP\_Group\_202\_201:

```

Group Type: External  AS: 202  Local AS: 201  Flags: <>
Total Peers: 1        Established Peers:

```

---

Another, more complete example of a BGP dump file follows:

---

```

          BGP.0.0.0.0+179 Priority 40      Port  179      Socket 11
RtProto
BGP
          Trace options: packets
          Trace file: /var/tmp/gated.log  size 1024000  files 3

          BGP_201.192.168.10.50+33073    Priority 50    Port  33073
Socket 1
2      RtProto BGP      RtBit 3
          Trace options: packets
          Trace file: /var/tmp/gated.log  size 1024000  files 3

          BGP_201.192.168.10.50+33073_Traffic
          last: 14:47:13.130      next: 14:47:13.130

          BGP_Group_201_202      Priority 50    RtProto BGP
          Trace options: packets
          Trace file: /var/tmp/gated.log  size 1024000  files 3

```

ASMatch Priority 50

ASPaths Priority 50

<.....>

```
223.239          mask 255.255.255
                  entries 1          announce 1
                  TSI:
                  Instability Histories:

                  *BGP Preference: 170          Source: 192.168.10.50
                  NextHop: 192.168.10.254          Interface:

192.168.10.32
(hme1)

                  State:
                  Local AS: 202 Peer AS: 201
                  Age: 10 Metric: -1          Metric2: -1          Tag: 0
                  Task: BGP_201.192.168.10.50+33073
                  Announcement bits(1): 2-KRT
                  AS Path: (202) {201}IGP (Id 6)

223.249          mask 255.255.255
                  entries 1          announce 1
                  TSI:
                  BGP 192.168.10.50 (External AS 201) no metrics
                  Instability Histories:

                  *Static Preference: 60
                  NextHop: 192.168.10.250          Interface:

192.168.10.32
(hme1)

                  State:
                  Age: 7 Metric: 0          Metric2: 0          Tag: 0
                  Task: RT
                  Announcement bits(2): 2-KRT 3-
BGP_201.192.168.10.50+3307
3
```

AS Path: IGP (Id 1)

<.....>

Task BGP.0.0.0.0+179:

Groups: 1 Peers: 1 (1 configured) Active Incoming: 0  
 Free Peers: 0 Free Groups: 0 State: Listening  
 AS Path: AS Path Regular Expression: (.\*)

<.....>

Task BGP\_201.192.168.10.50+33073:

Peer: 192.168.10.50+33073 Local: 192.168.10.32+179 Type:  
 External

State: Established Flags: <>

Last State: OpenConfirm Last Event: RecvKeepAlive Last Error:  
 None

Options: <>

Peer Version: 4 Peer ID: 0.0.0.50 Local ID: 0.0.0.5  
 Active Holdtime: 180

Last traffic (seconds): Received 10 Sent 22 Checked 27

Input messages: Total 7 Updates 5 Octets 270

Output messages: Total 7 Updates 4 Octets 244

Route Queue Timer: unset Route Queue: empty

Task BGP\_Group\_201\_202:

Group Type: External AS: 201 Local AS: 202 Flags: <>

Total Peers: 1 Established Peers: 1

Task ASMatch:

Task ASPaths:

Id 1 Refs 7 Hash 0 Lengths: Path 0, Seg 0, Attr 0, Aggr 0, Alloc'd 96  
 Path: IGP

Id 2 Refs 1 Hash 75 Lengths: Path 2, Seg 2, Attr 0, Aggr 0, Alloc'd 96  
 Path: (202) 201 IGP

Id 4 Refs 1 Hash 75 Lengths: Path 2, Seg 2, Attr 0, Aggr 0, Alloc'd 96  
 Path: (202) 201 IGP

Id 6 Refs 1 Hash 75 Lengths: Path 2, Seg 2, Attr 0, Aggr 0, Alloc'd 96

Path: (202) 201 IGP

---

## 8.5.9 AS Paths

The next section of the dump will include AS path information for each ID, such as:

- references (for ID1, 18 references exist)
- hash (0)
- lengths (Path 0, Seg 0, Attr 0, Aggr 0, Alloc'd 100)
- path (IGP)

---

Task ASPaths:

Id 1Refs 18 Hash 0 Lengths: Path 0, Seg 0, Attr 0, Aggr 0, Alloc'd 100  
Path: IGP

---

## 8.5.10 Kernel

The final section of the dump will include kernel information, such as:

- options: <>
- kernel support: <reject blackhole varmask host>
- remnant timeout: 3:00
- routes in kernel: 18
- limit: unlimited
- flash install limit: 20
- flash install routes: interface
- background install limit: 120
- background install routes: low
- kernel install state: normal
- interface add/delete/change queue: 0
- queued High priority deletion queue: 0
- queued High priority change queue: 0
- queued High priority add queue: 0
- queued Normal priority deletion queue: 0
- queued Normal priority change queue: 0
- queued Normal priority add queue: 0 queued

---

Task KRT:

Options: <>      Kernel support: <reject blackhole varmask host>  
Remnant timeout: 3:00  
Routes in kernel: 18      Limit: unlimited  
Flash install limit: 20      Flash install routes: interface  
Background install limit: 120      Background install routes: low

Kernel install state: normal

Interface add/delete/change queue: 0 queued

High priority deletion queue: 0 queued

High priority change queue: 0 queued

High priority add queue: 0 queued

Normal priority deletion queue: 0 queued

Normal priority change queue: 0 queued

Normal priority add queue: 0 queued

---

## 8.6 Dump vs. GII

The GII and dump facilities both have route table and protocol information. The dump is a more detailed listing of the internal state, while GII tends to be easier to read for quick status information. The GII interface can also be easier to use for those familiar with other vendors' command-line interfaces. The dump file should be emailed to NextHop Technologies when requesting information regarding operational issues or problems.

