

NAME

racluster – aggregate **argus(8)** data files.

SYNOPSIS

racluster [**-f** *conf*] [**-m** *agr(s)*] [**-M** *mode(s)*] [**raoptions**]

DESCRIPTION

Racluster reads **argus** data from an *argus-data* source, and clusters/merges the records based on the flow key criteria specified either on the command line, or in a **racluster** configuration file, and outputs a valid *argus-stream*. This tool is primarily used for data mining, data management and report generation.

The default action is to merge status records from the same flow and argus probe, providing in some cases huge data reduction with limited loss of flow information. **Racluster** provides the ability to modify the flow model key, either using the "-m" option, or in the **racluster.conf** file, allowing records to be clustered based on any number of attributes. This supports the development of important reports, such as MPLS LSP usage statistics, DiffServe flow marking policy verification, VLAN group behavior, IP distance related measurements, routing loop detection, traceroute path data recovery, and complex availability/reachability reports, to name just a few useful applications.

Please see **racluster.5** for detailed information regarding **racluster** configuration.

OPTIONS

Racluster, like all **ra** based clients, supports a number of **ra options** including filtering of input argus records through a terminating filter expression, and the ability to specify the output style, format and contents for printing data. See **ra(1)** for a complete description of **ra options**. **racluster(1)** specific options are:

-m *aggregation object*

Supported aggregation objects are:

none	use a null flow key.
srcid	argus source identifier.
smac	source mac(ether) addr.
dmac	destination mac(ether) addr.
smpls	source mpls label.
dmpls	destination label addr.
svlan	source vlan label.
dvlan	destination vlan addr.
saddr/[l m]	source IP addr/[cidr len m.a.s.k].
daddr/[l m]	destination IP addr/[cidr len m.a.s.k].
matrix/l	sorted src and dst IP addr/cidr len.
proto	transaction protocol.
sport	source port number.
dport	destination port number.
stos	source TOS byte value.
dtos	destination TOS byte value.
sttl	src -> dst TTL value.
dttl	dst -> src TTL value.
stcpb	src -> dst TCP base sequence number.
dtcpb	dst -> src TCP base sequence number.
inode/[l m]	intermediate node IP addr/[cidr len m.a.s.k], source of ICMP mapped events.
sco	source ARIN country code, if present.
dco	destination ARIN country code, if present.
sas	source node origin AS number, if available.
das	destination node origin AS number, if available.
ias	intermediate node origin AS number, if available.

-M *modes*

Supported modes are:

correct	Attempt to correct the direction of flows by also searching the reverse flow key, if a match isn't found in the cache. This mode is on by default when using the default full 5-tuple flow key definitions.
nocorrect	Turn off flow correction for direction. This mode is used by default if the flow key has been changed.
norep	Do not generate an aggregate statistic for each flow. This is used primarily when the output represents a single object. Primarily used when merging status records to generate single flows that represent single transactions.
rmon	Generate data suitable for producing RMON types of metrics.
ind	Process each input file independantly, so that after the end of each inputfile, racluster flushes its output.
replace	Replace each inputfile contents, with the aggregated output.

-V Verbose operation, printing a line of output for each input file processed. Very useful when using the ra() -R option.

INVOCATION

A sample invocation of **racluster(1)**. This call reads **argus(8)** data from **inputfile** and aggregates the IP protocol based **argus(8)** data, based on the source and destination address matrix and the destination port used by tcp flows, and report the metrics as a percent of the total. For most services, this provides service specific metrics on a client/server basis.

```
racluster -% -r inputfile -m saddr daddr dport - \
    tcp and syn and synack
```

This call reads **argus(8)** data from **inputfile** and generates the path information that traceroute use would generate (assuming that traceroute was run during the observation period).

```
racluster -r inputfile -m saddr daddr sttl inode -w - - icmpmap | \
rasort -m sttl -s saddr dir daddr inode avgdur spkts
```

COPYRIGHT

Copyright (c) 2000-2011 QoSient. All rights reserved.

SEE ALSO

racluster(5), **ra(1)**, **rarc(5)**, **argus(8)**,

FILES**AUTHORS**

Carter Bullard (carter@qosient.com).

BUGS