SNMP Trace Analysis

Jürgen Schönwälder

IM 2007, Munich, 2007-05-23

Co-authors: A. Pras, M. Harvan, J. Schippers, R. van de Meent
Support: EU IST-EMANICS Network of Excellence (#26854), SURFnet and various network operators

Outline

Outline

Who wants to be an SNMP millionaire?

Who wants to be an SNMP millionaire?

Question

What is the most widely used version of SNMP?

Answers

a) SNMPv1

b) SNMPv2c

c) SNMPv3

d) SNMPv4

Who wants to be an SNMP millionaire?

Question

What is the most widely used version of SNMP?

Answers

a) SNMPv1

b) SNMPv2c

c) SNMPv3

d) SNMPv4

Who wants to be an SNMP millionaire?

Question

What is a commonly used max-repetitions parameter value?

Answers

a) 1

b) 2

c) 42

d) 1000

Who wants to be an SNMP millionaire?

Question

What is a commonly used max-repetitions parameter value?

Answers

a) 1

b) 2

c) 42

d) 1000

Who wants to be an SNMP millionaire?

Question

How many agents does an SNMP manager typically manage?

Answers

a) 2 agents

b) 42 agents

c) 288 agents

d) >288 agents

Who wants to be an SNMP millionaire?

Question

How many agents does an SNMP manager typically manage?

Answers

a) 2 agents

b) 42 agents

c) 288 agents

d) >288 agents

Who wants to be an SNMP millionaire?

Question

What is the average size of an SNMP message?

Answers

a) 42 bytes

c) 767 bytes

b) 484 bytes

d) 1500 bytes

Who wants to be an SNMP millionaire?

Question

What is the average size of an SNMP message?

Answers

a) 42 bytes

b) 484 bytes

c) 767 bytes

d) 1500 bytes

We all know SNMP...

- The Simple Network Management Protocol (SNMP) is widely deployed to
  - monitor devices (collect statistics, event reports),
  - control devices (turning knobs), and
  - (to a lesser extent) configure devices
- SNMP supports "fancy" features to allow applications to do the right thing
  - discontinuity indicators
  - row creation modes
  - advisory locks
  - ...
- SNMP technology is well documented and understood (if you care to study the right documents)

. . . but not how it is used in practice!

- What are typical SNMP usage patterns?
- Which table retrieval algorithms are popular?
- How to model the arrival process of SNMP messages?
- What is the size distribution of SNMP messages?
- Which features of SNMP are used/not used?
- Which MIB objects (MIB modules) are frequently used?
- Is trap-directed polling reality or a myth?
- Are the fully automated control loops?
- Are SNMP improvements relevant for deployments?
- . . .

So why is this important?

1. Researchers write papers how to improve SNMP or how other technologies (e.g., Web Services) compare to SNMP without having a justified model
2. The IETF works on extensions (e.g., session-based security in ISMS) without knowing network management traffic models (and in the context of ISMS to what extend a session-based approach to security is viable)
3. The IETF requires features during MIB design/review without knowing whether they are used in practice
4. Implementors always want to know which features are worth to spend development time on

So why is this important?

1. Researchers write papers how to improve SNMP or how other technologies (e.g., Web Services) compare to SNMP without having a justified model
2. The IETF works on extensions (e.g., session-based security in ISMS) without knowing network management traffic models (and in the context of ISMS to what extend a session-based approach to security is viable)
3. The IETF requires features during MIB design/review without knowing whether they are used in practice
4. Implementors always want to know which features are worth to spend development time on

So why is this important?

1. Researchers write papers how to improve SNMP or how other technologies (e.g., Web Services) compare to SNMP without having a justified model
2. The IETF works on extensions (e.g., session-based security in ISMS) without knowing network management traffic models (and in the context of ISMS to what extend a session-based approach to security is viable)
3. The IETF requires features during MIB design/review without knowing whether they are used in practice
4. Implementors always want to know which features are worth to spend development time on

So why is this important?

1. Researchers write papers how to improve SNMP or how other technologies (e.g., Web Services) compare to SNMP without having a justified model

2. The IETF works on extensions (e.g., session-based security in ISMS) without knowing network management traffic models (and in the context of ISMS to what extend a session-based approach to security is viable)

3. The IETF requires features during MIB design/review without knowing whether they are used in practice

4. Implementors always want to know which features are worth to spend development time on

Outline

Measurements

- Goal: Capture SNMP traffic from operational networks to model how network management protocols are used in practice

- Measurement process:
  1. Capture raw SNMP traces in `pcap` capture files
  2. Convert raw traces into an intermediate format
  3. Filter traces to suppress / anonymize sensitive data
  4. Store filtered / anonymized traces in a repository
  5. Run analysis scripts on filtered / anonymized traces

- Intermediate formats [1]:
  - `xml` - sexy and comprehensive but expensive
  - `csv` - classic and efficient but restricted

Traces contain sensitive data

- Dealing with sensitive data:
  - In some cases, the operator chooses to provide raw traces to researchers under an NDA
  - In some cases, the operator chooses to provide filtered / anonymized traces to researchers under an NDA
  - In some cases, the operator chooses to keep traces under local control and commits to run analysis scripts on them and to provide the results
- Our experience / recommendation:
  - Tools should support many different approaches
  - Trust between researchers and operators is key
  - Building trust relationships is a good investment

Tool support (snmpdump)

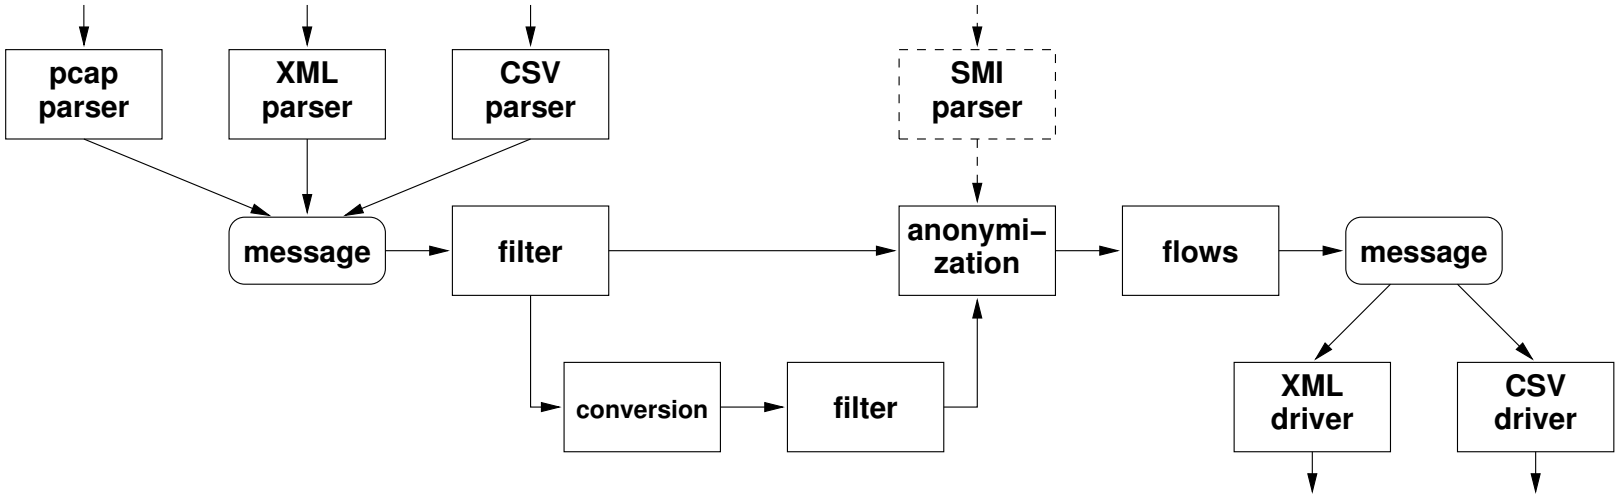

- snmpdump reads as input pcap or xml or csv files and produces output in xml or csv format
- snmpdump takes care of datagram reassembly (libnids)
- snmpdump can split a trace into flows (see below)

Filtering, conversion, anonymization

- The filter module is responsible to remove information that should not be made available (filter-out principle)
  - Filtering must happen as early as possible
  - A filter is a regex matched against field names
  - Typically used to remove community strings
- The conversion module implements trap conversion as specified in RFC 3584 [2] section 3.1
  - Single trap format simplifies scripts
  - Filtering has to be re-applied after conversion
  - Conversion is optional
- The anonymization module scrambles data in order to raise the effort needed to obtain sensitive information about the internals of an operational network (see [3])

Precise and effective to compute metrics needed...

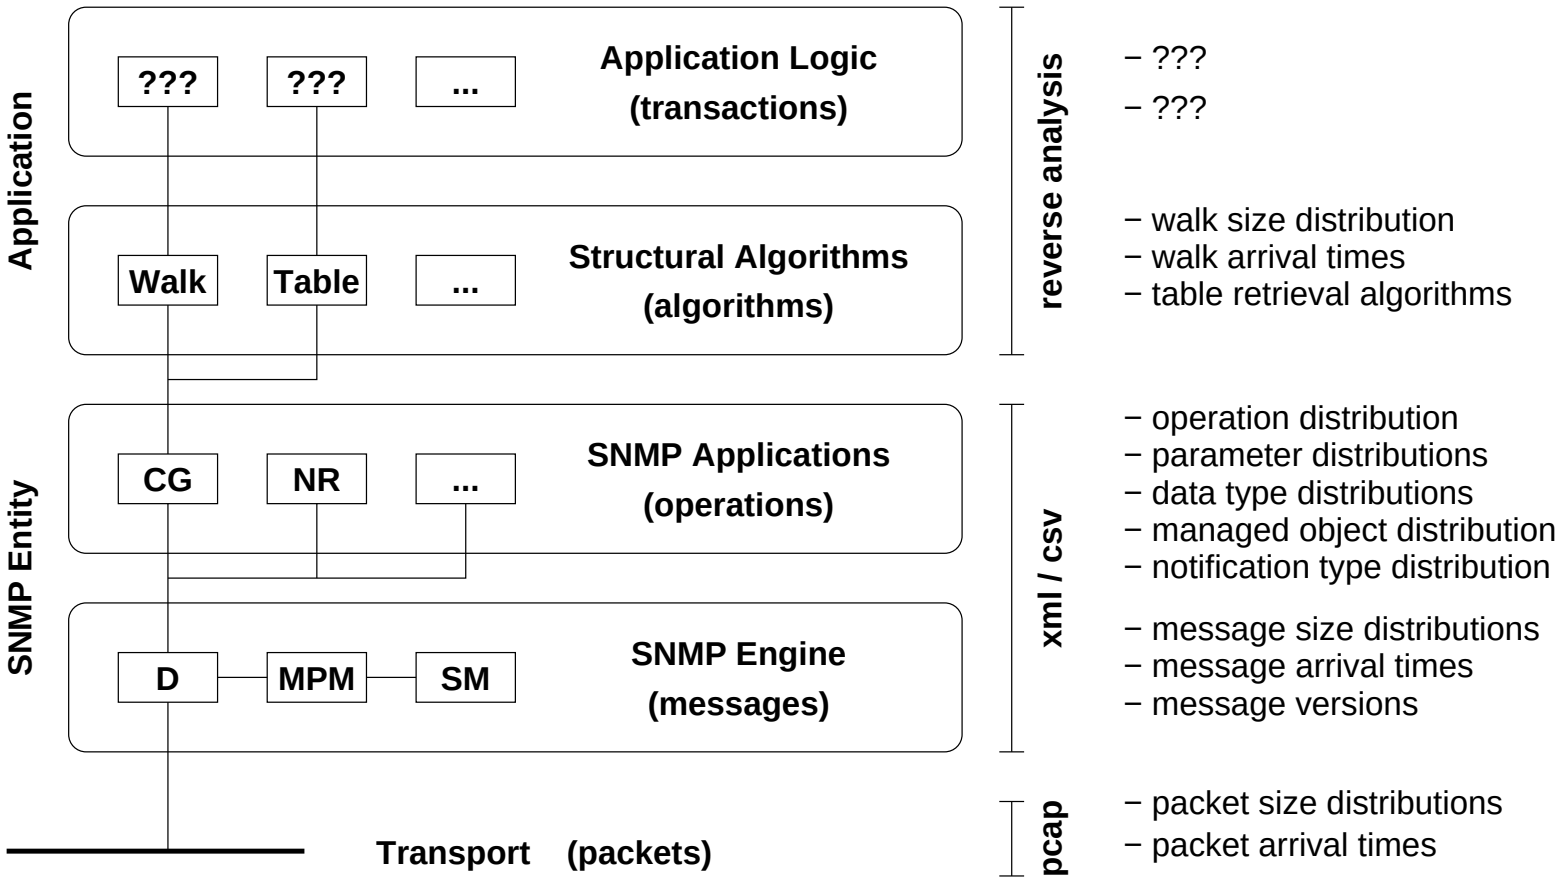

- Layered model of metrics to abstract from the concrete details of the management protocols used

Outline

Trace collection

trace	description	start	hours
l01t02	national research network	2005-07-26	162.98
l01t05	national research network	2006-07-10	336.00
l02t01	university network	2006-04-21	294.62
l03t02	faculty network	2006-04-27	159.21
l04t01	server-hosting provider	2006-04-14	4.00
l05t01	regional network provider	2006-04-19	580.60
l06t01	national research network	2006-05-14	222.08
l12t01	point of presence	2006-07-10	208.02

- Additional traces are being collected within EMANICS
- More traces are always welcome!!

Trace characterization

trace	size [MB]	messages	SNMPv1	SNMPv2	SNMPv3
l01t02	6369	51772136	100.0%	-	-
l01t05	14043	40072529	-	100.0%	0.0%
l02t01	77789	258010521	5.5%	94.5%	-
l03t02	130858	871361365	95.0%	5.0%	-
l04t01	10	15099	35.7%	64.3%	-
l05t01	2898	25298667	100.0%	-	-
l06t01	24683	89277889	57.4%	42.6%	-
l12t01	312	2619884	32.3%	67.7%	-

- One trace included a few very sporadic SNMPv3 packets (someone testing SNMPv3?)
- Organization *l01* uses in one location 100% SNMPv1 and in a second location 100% SNMPv2c

Protocol operations

trace	Get	Next	Bulk	Set	Trap	Inform	Resp
l01t02	0.0	50.0	-	0.0	-	-	50.0
l01t05	0.0	-	50.0	-	-	-	50.0
l02t01	0.1	2.4	47.1	0.0	0.7	-	49.6
l03t02	0.3	49.8	-	0.0	0.0	-	49.9
l04t01	32.8	3.8	22.9	-	-	-	40.5
l05t01	50.0	0.0	-	-	0.0	-	50.0
l06t01	12.1	31.4	6.5	-	0.0	0.0	50.0
l12t01	1.0	49.0	-	-	0.0	-	49.9

- In *l01t02*, all `Set` operations were trying to modify `sysLocation` with a value of type `Integer32`
- In *l02t01* and *l03t02*, `Set` operations were used to trigger download of config information
- In *l04t01*, there were significantly more requests than responses (system maintenance)

Protocol operation parameters

trace	Get	Next	Bulk	max-reps	non-reps
l01t02	37.5%	99.3%	-	-	-
l01t05	100.0%	-	100.0%	10/50	0
l02t01	56.3%	99.9%	100.0%	1/10/20/25	0
l03t02	1.6%	99.9%	-	-	-
l04t01	100.0%	100.0%	100.0%	1000	0
l05t01	99.9%	95.6%	-	-	-
l06t01	8.7%	2.6%	0.0%	12	0
l12t01	100.0%	99.9%	-	-	-

- Except for *l06t01*, single varbind `GetNext` and `GetBulk` operations dominate
- In *l06t01*, 86.5% of the `GetBulk` operations contain two varbinds and the remaining 13.5% contain eight varbinds

Response size distribution

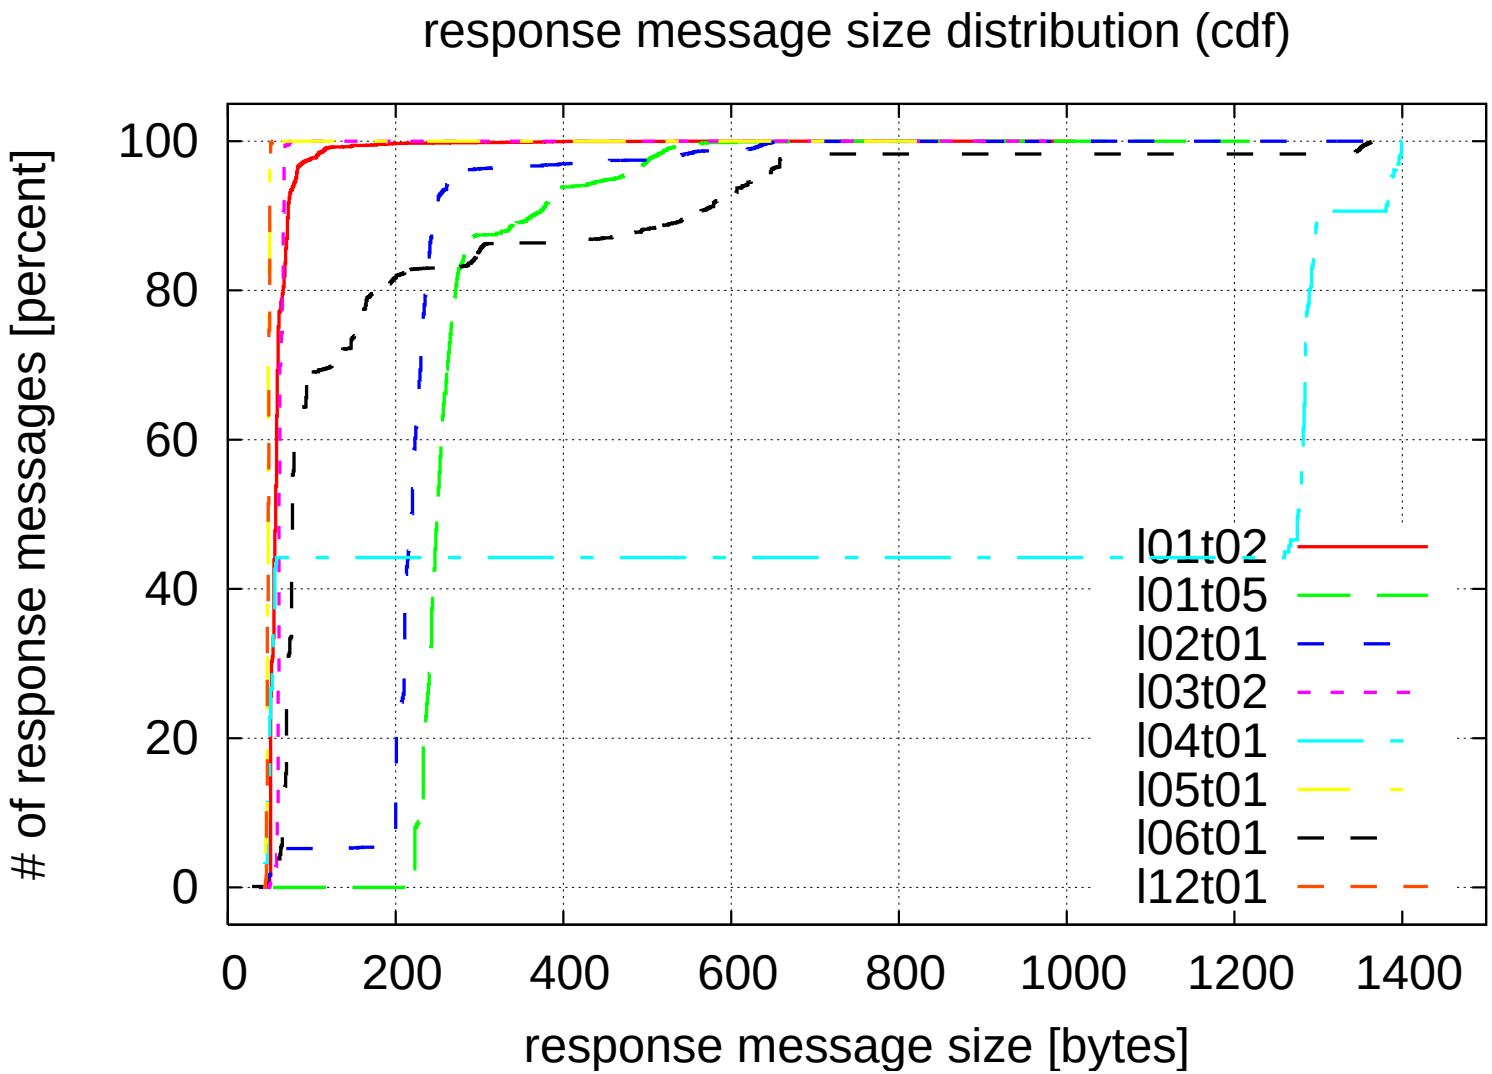

SNMP message flows

Definition

An SNMP message flow is defined as all messages between a source and destination address pair which belong to a command generator (CG) / command responder (CR) relationship or a notification originator (NO) / notification receiver (NR) relationship.

- The above definition deliberately does not consider port numbers (they change too frequently)
- Multi-homed managers will appear with multiple flows

Flow statistics

trace	cg/cr flows	no/nr flows	cg	cr	no	nr
l01t02	203	-	3	178	-	-
l01t05	8	-	2	8	-	-
l02t01	258	197	5	240	197	1
l03t02	42	20	25	20	17	2
l04t01	34	-	3	34	-	-
l05t01	117	2	9	99	2	2
l06t01	288	125	3	260	125	2
l12t01	30	6	5	26	6	1

- Traffic is not evenly distributed across the flows
- Most traces have very few dominating flows

Flow size distribution

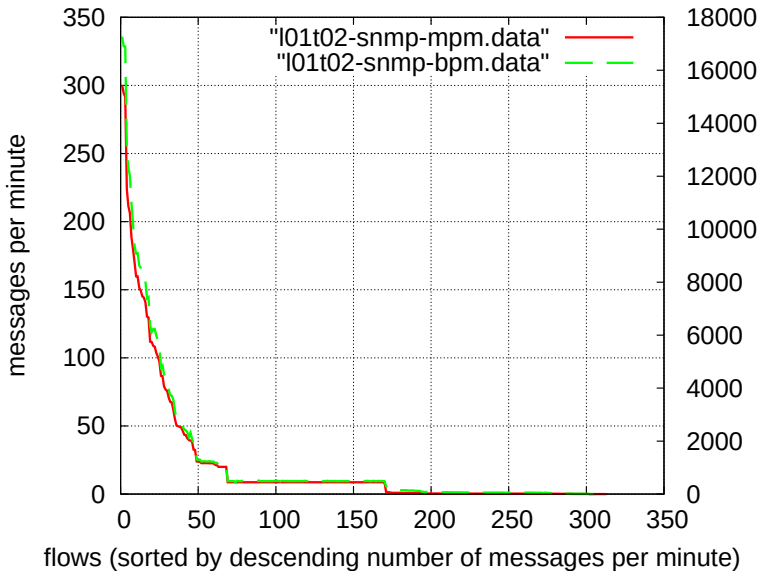

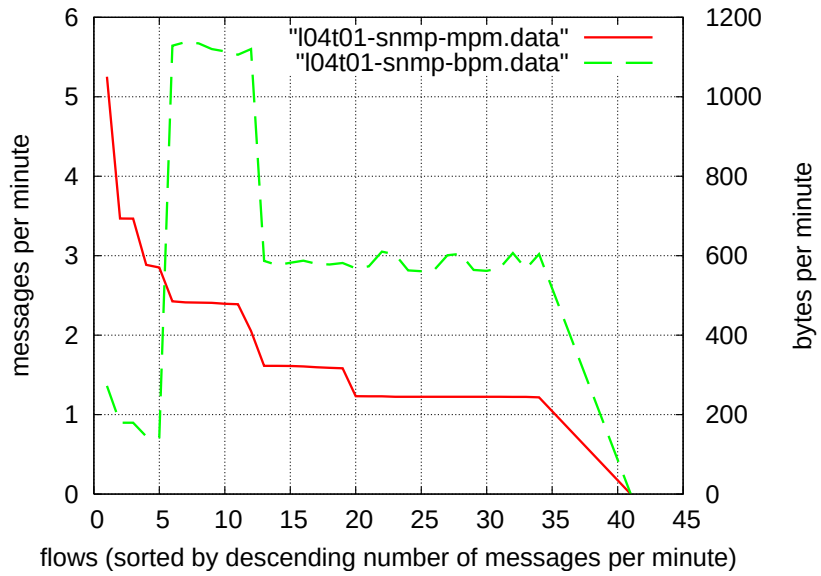

- In some traces, the number of message and bytes per flow is closely correlated
- In other traces, this is not the case (essentially due to `GetBulk` usage)

Flow topology *l04t01*

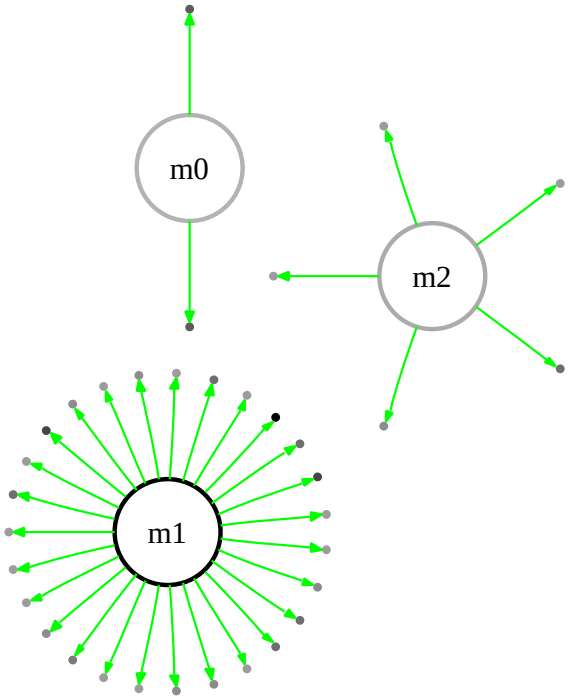

- Typical simple monitoring topology
- Some flows carry more traffic than other flows
- Gray-level indicates intensity (but difficult to see)

Flow topology *l01t02*

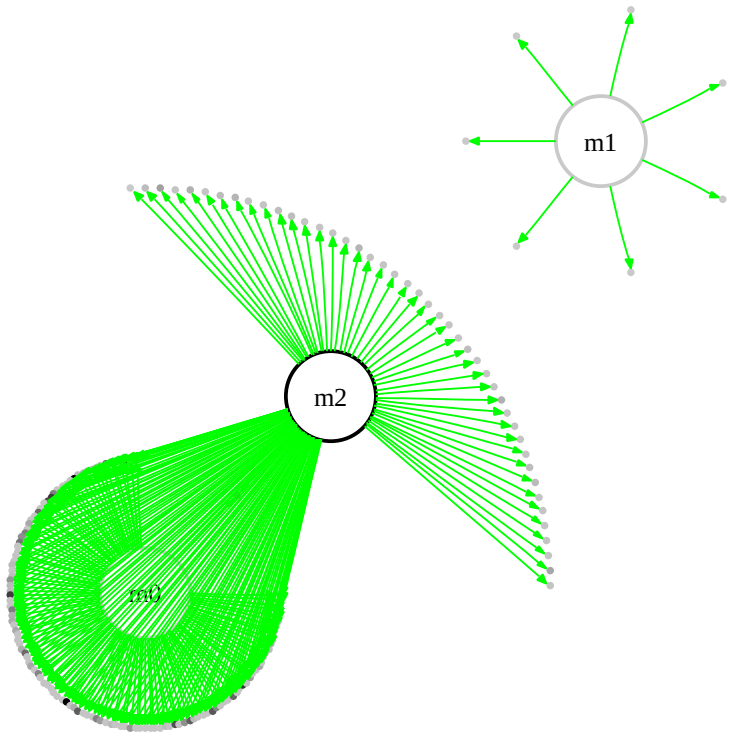

- Slightly more complex monitoring topology
- Some devices are interacting with multiple management interfaces

Flow topology *l03t02*

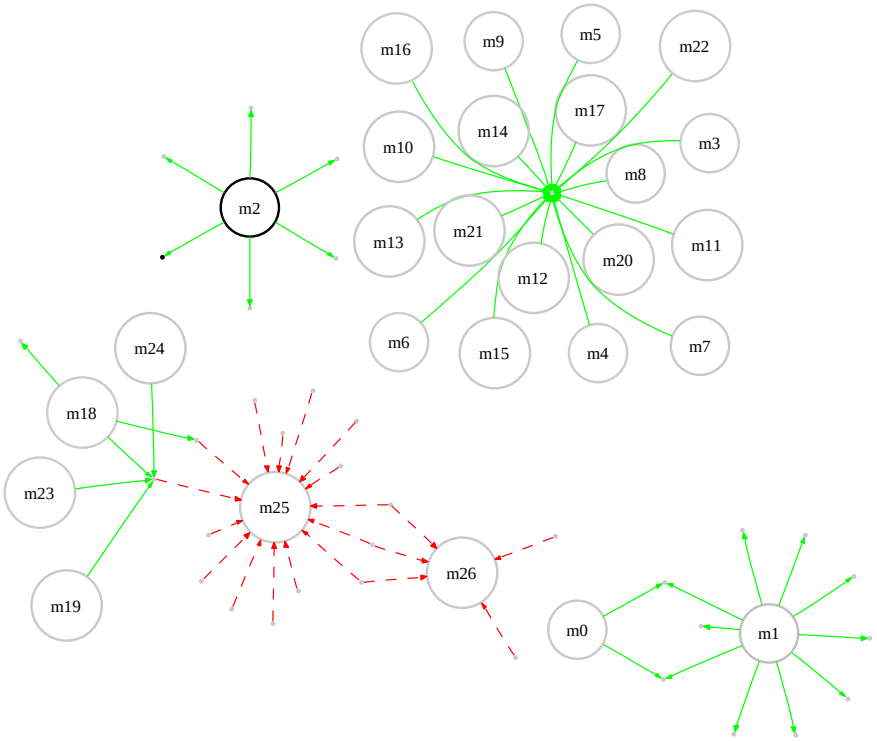

- Large number of managers talking to a single agent
- Analysis revealed that the device is a printer queried from a PC which is getting dynamically changing IP addresses

Flow topology *l06t01*

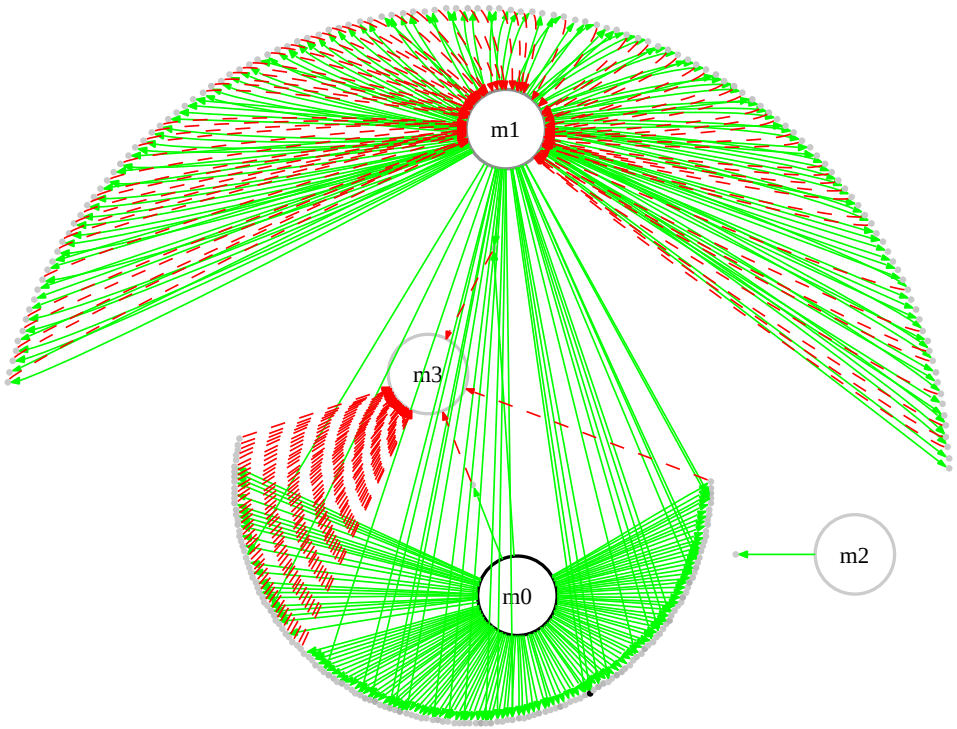

- Significant number of NO/NR flows
- Single pretty dark dot indicates that a single flow is dominating

Flow topology *l06t01* (treemap)

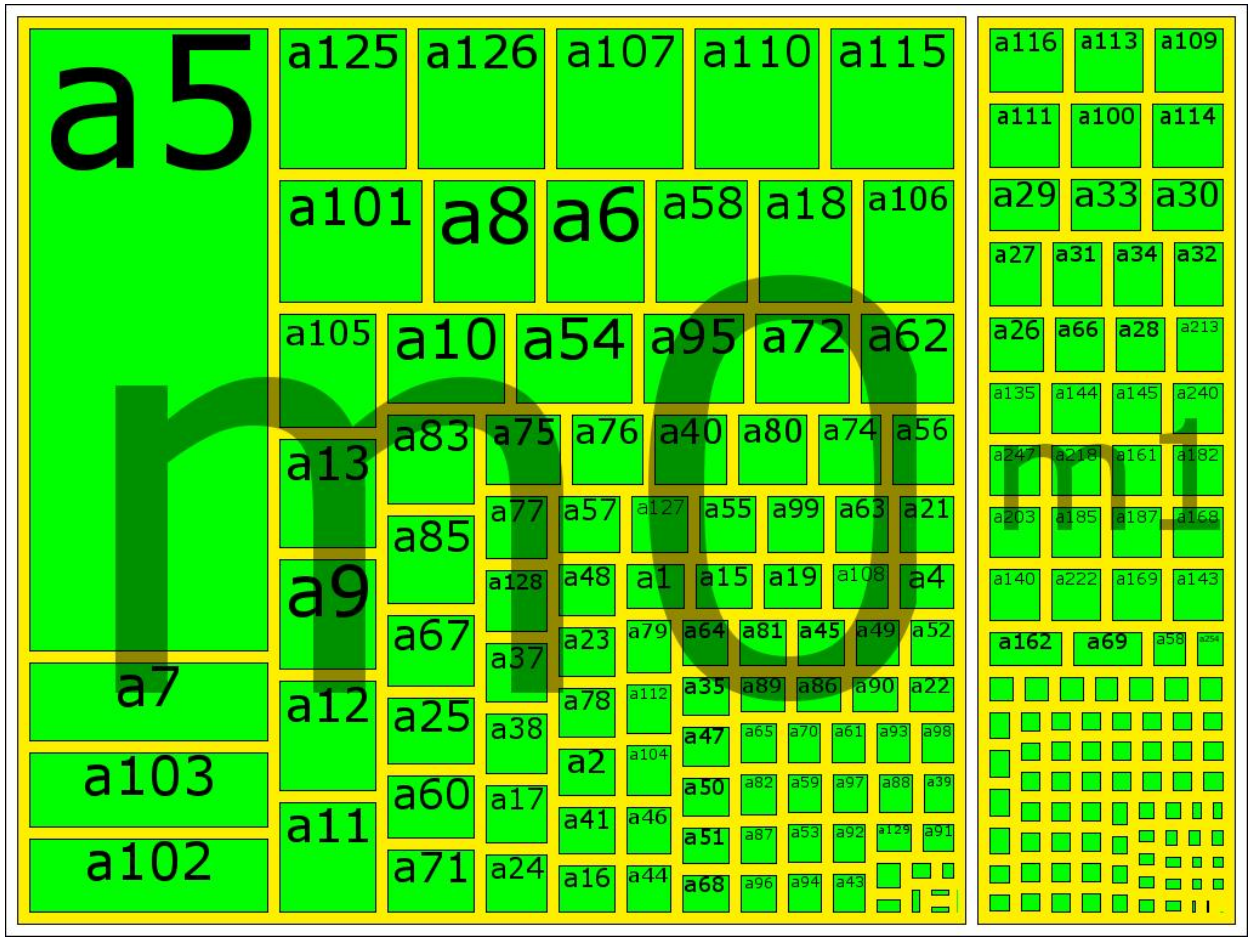

- Treemap plots nicely visualizes contribution of the flows

Data type usage

trace	int32	uint32	uint64	oct	oid	ip	null	exc
l01t02	48.1	3.2	-	39.6	0.3	8.6	0.2	-
l01t05	13.4	21.0	52.7	12.9	0.0	0.0	-	0.0
l02t01	22.4	45.1	18.4	11.7	2.4	0.0	0.0	0.0
l03t02	2.5	95.0	-	2.4	0.1	0.0	0.0	-
l04t01	0.7	0.5	98.8	-	-	-	-	-
l05t01	2.6	80.1	-	17.0	0.0	0.0	-	-
l06t01	37.9	23.8	7.5	30.7	0.0	0.0	0.0	0.0
l12t01	48.3	51.5	0.0	0.1	0.1	0.0	0.0	-

- Strong dominance of integral types; some traces contain in addition a significant portion of octet string data
- Some read-only string objects (e.g., `ifDescr`) are retrieved over and over again

Managed objects usage

trace	IF	BR	BGP	HR	ENT	CIS	CP
l01t02	40.1	-	17.6	-	10.3	30.4	-
l01t05	99.7	-	-	0.0	-	-	-
l02t01	93.5	5.5	0.0	-	0.1	0.0	-
l03t02	33.3	65.1	0.0	0.0	0.0	0.1	-
l04t01	99.7	-	-	-	-	-	-
l05t01	80.1	-	-	-	-	-	17.0
l06t01	91.3	0.0	0.0	-	0.0	2.0	-
l12t01	50.4	0.0	-	47.7	-	-	-

- In trace *l06t01*, 32-bit counters dominate and the discontinuity indicator is ignored
- In trace *l02t01*, *all* columns of the `ifTable` and the `ifXTable` are retrieved regularly

Notifications

- In trace *l02t01*, about 52.1% of the notifications are fan failure notifications that are repeated periodically. Some 42.2% are interface up/down notifications while the remaining notifications are HP and Avaya specific
- In trace *l03t02*, we found that all notifications were reporting printer problems
- Trace *l05t01* contains only Cisco notifications related to TCP session teardowns and configuration changes
- Trace *l06t01* has 26.1% BGP and 8.1% PIM routing related notifications. Some 20.0% are sensor threshold crossing notifications while 13.2% are Cisco notifications related to TCP session teardowns
- Trace *l12t01* contains 68.5% authentication failure, 14.8% cold start and 16.7% shut down notifications

Notifications

- In trace *l02t01*, about 52.1% of the notifications are fan failure notifications that are repeated periodically. Some 42.2% are interface up/down notifications while the remaining notifications are HP and Avaya specific
- In trace *l03t02*, we found that all notifications were reporting printer problems
- Trace *l05t01* contains only Cisco notifications related to TCP session teardowns and configuration changes
- Trace *l06t01* has 26.1% BGP and 8.1% PIM routing related notifications. Some 20.0% are sensor threshold crossing notifications while 13.2% are Cisco notifications related to TCP session teardowns
- Trace *l12t01* contains 68.5% authentication failure, 14.8% cold start and 16.7% shut down notifications

Notifications

- In trace *l02t01*, about 52.1% of the notifications are fan failure notifications that are repeated periodically. Some 42.2% are interface up/down notifications while the remaining notifications are HP and Avaya specific
- In trace *l03t02*, we found that all notifications were reporting printer problems
- Trace *l05t01* contains only Cisco notifications related to TCP session teardowns and configuration changes
- Trace *l06t01* has 26.1% BGP and 8.1% PIM routing related notifications. Some 20.0% are sensor threshold crossing notifications while 13.2% are Cisco notifications related to TCP session teardowns
- Trace *l12t01* contains 68.5% authentication failure, 14.8% cold start and 16.7% shut down notifications

Notifications

- In trace *l02t01*, about 52.1% of the notifications are fan failure notifications that are repeated periodically. Some 42.2% are interface up/down notifications while the remaining notifications are HP and Avaya specific
- In trace *l03t02*, we found that all notifications were reporting printer problems
- Trace *l05t01* contains only Cisco notifications related to TCP session teardowns and configuration changes
- Trace *l06t01* has 26.1% BGP and 8.1% PIM routing related notifications. Some 20.0% are sensor threshold crossing notifications while 13.2% are Cisco notifications related to TCP session teardowns
- Trace *l12t01* contains 68.5% authentication failure, 14.8% cold start and 16.7% shut down notifications

Notifications

- In trace *l02t01*, about 52.1% of the notifications are fan failure notifications that are repeated periodically. Some 42.2% are interface up/down notifications while the remaining notifications are HP and Avaya specific
- In trace *l03t02*, we found that all notifications were reporting printer problems
- Trace *l05t01* contains only Cisco notifications related to TCP session teardowns and configuration changes
- Trace *l06t01* has 26.1% BGP and 8.1% PIM routing related notifications. Some 20.0% are sensor threshold crossing notifications while 13.2% are Cisco notifications related to TCP session teardowns
- Trace *l12t01* contains 68.5% authentication failure, 14.8% cold start and 16.7% shut down notifications

Outline

This is just the beginning...

- Started with SNMP because
  - it is widely deployed and we understand it well
  - it is reasonably complex to start with
- Work is underway
  - to analyze periodic / aperiodic behaviour
  - to investigate walks / table retrieval algorithms
  - to identify semantic management transactions
- Plans exist to cover additional protocols
  - SYSLOG, CLI, RADIUS, NETFLOW, ...
- More traces will be added
  - to cover more network types (campus, enterprise, ISP)
  - to increase statistical evidence in our results

. . . of a longer journey.

Long term goals

- Develop models of management plane interactions
- Abstracting away from protocol specifics
- Definition of management plane metrics (IPPM like)
- Simulation models to evaluate approaches / scenarios
- Advanced visualization techniques to explore data sets
- Understand what scalability actually means

The end

Questions?

. . . of a longer journey.

Long term goals

- Develop models of management plane interactions
- Abstracting away from protocol specifics
- Definition of management plane metrics (IPPM like)
- Simulation models to evaluate approaches / scenarios
- Advanced visualization techniques to explore data sets
- Understand what scalability actually means

The end

Questions?

References

J. Schönwälder.

SNMP Traffic Measurements.
Internet Draft (work in progress) <draft-schoenw-nmrg-snmp-measure-01.txt>, International University Bremen, February 2006.

R. Frye, D. Levi, S. Routhier, and B. Wijnen.

Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework.
RFC 3584, Vibrant Solutions, Nortel Networks, Wind River Systems, Lucent Technologies, August 2003.

M. Harvan and J. Schönwälder.

Prefix- and Lexicographical-order-preserving IP Address Anonymization.
In *10th IEEE/IFIP Network Operations and Management Symposium*, pages 519–526, April 2006.

J. Schönwälder.

Characterization of SNMP MIB Modules.
In *Proc. 9th IFIP/IEEE International Symposium on Integrated Network Management*, pages 615–628. IEEE, May 2005.

J. Schönwälder, A. Pras, M. Harvan, J. Schippers, and R. van de Meent.

SNMP Traffic Analysis: Approaches, Tools, and First Results.
In *Proc. 10th IFIP/IEEE International Symposium on Integrated Network Management*, May 2007.