A Few Months In The Life Of An RPKI Validator

Rob Austein <sra@hactrn.net>
Randy Bush <randy@psg.com>
Michael Elkins <Michael.Elkins@sparta.com>
. . . and a lot of help from our friends

IEPG
Paris
25 March 2012

Introduction
Performance Graphs
Object Counts
Connection Counts
Objects/Connection
Seconds/Object
Average Connection Duration
Failure Rate
Rate Limiting
Repository Summaries
Conclusion

The World As Seen By One RPKI Validator

- Data as logged by one validator in Seattle.
- Data collection started late October 2011.
- Guilty parties are good people, all friends here.
- Expect updated report(s) at later date(s).

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

A Brief Overview of RPKI Validation

- Distributed global database of X.509 certificates and dependent objects.
- The X.509 certificates contain `rsync://` URIs.
- Validation starts at trust anchor(s).
- Validator walks certificate tree, following URIs.
- rcynic is one such validator.
- rcynic is session-oriented (cron job).

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Object Counts (Linear)

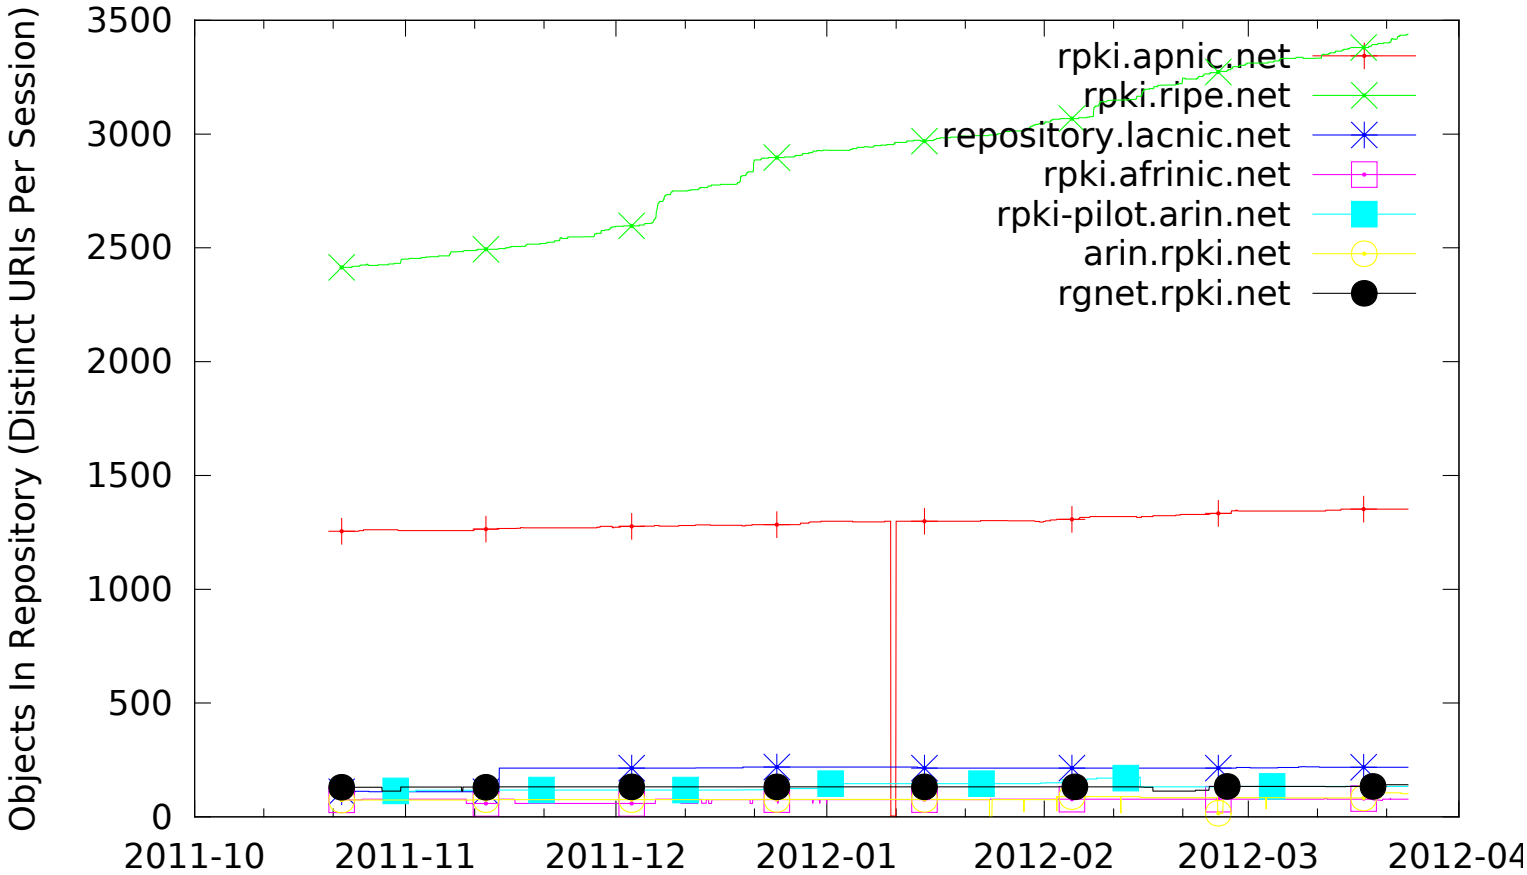

Object Counts (Logarithmic)

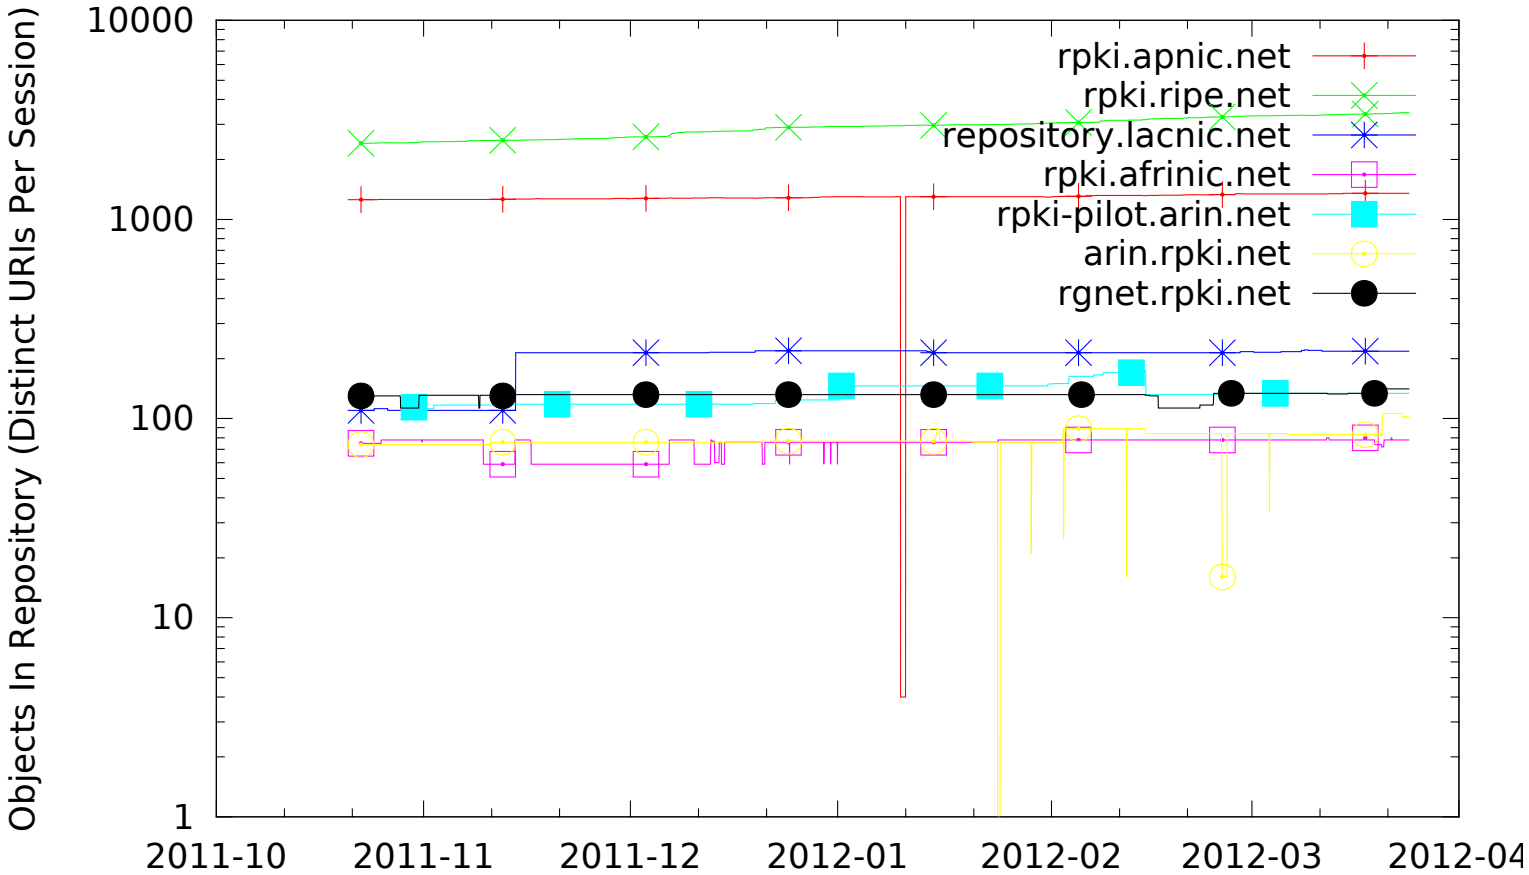

Object Counts: Observations

- Large downward spikes are either genuine mass extinction events or, more likely, validation failure of a high-level certificate causing a large subtree to go invalid. Either way, these usually indicate Something Very Bad.

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Connection Counts (Linear)

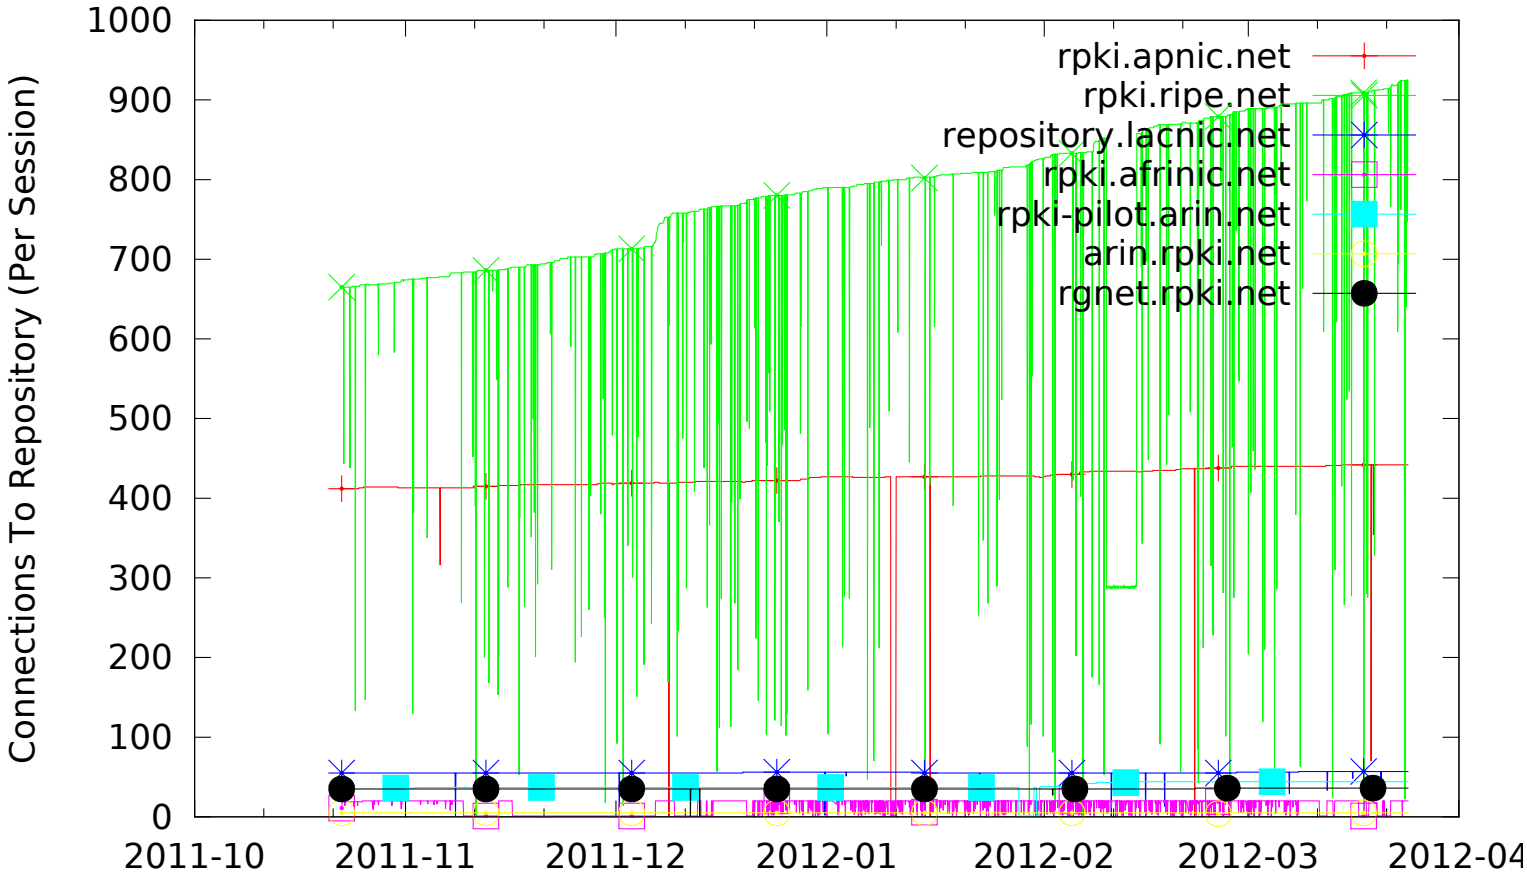

Connection Counts (Logarithmic)

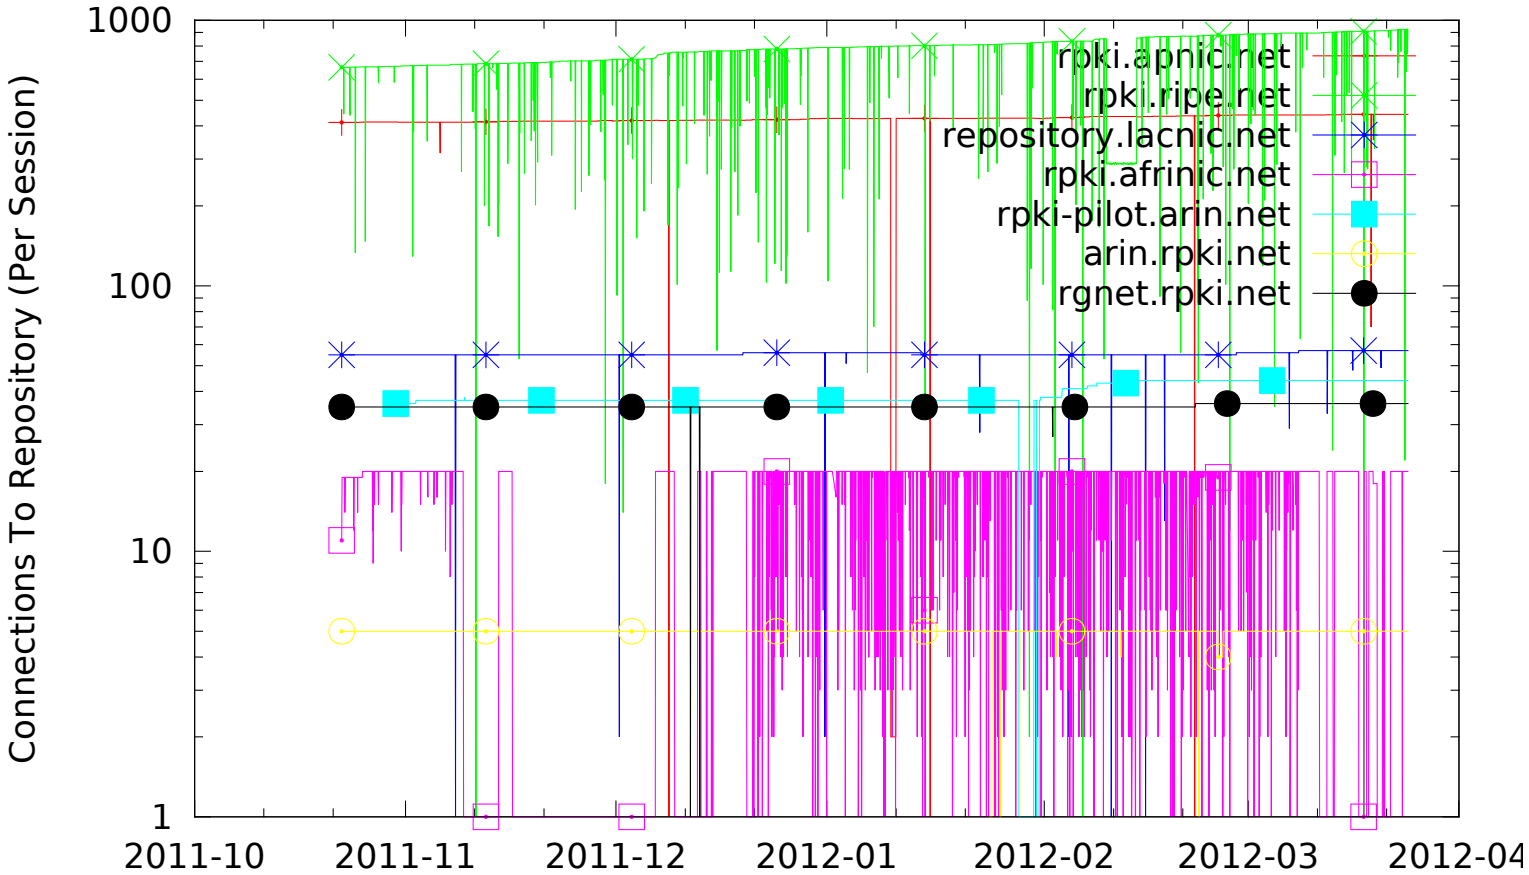

Connection Counts: Observations

- Downward spikes are connection failures, because once we decide a repository server is down, we give up on it until the next session.
- Are those repositories really that flaky? Perhaps, but at least one of them does their own monitoring and says not. Problem only seems to occur for repositories with AAAA RRs. Uh oh. As far as we can tell this is an IPv6 problem: IPv6 from Seattle to Amsterdam appears to be much flakier than IPv4 from Seattle to Brisbane.

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Objects/Connection (Linear)

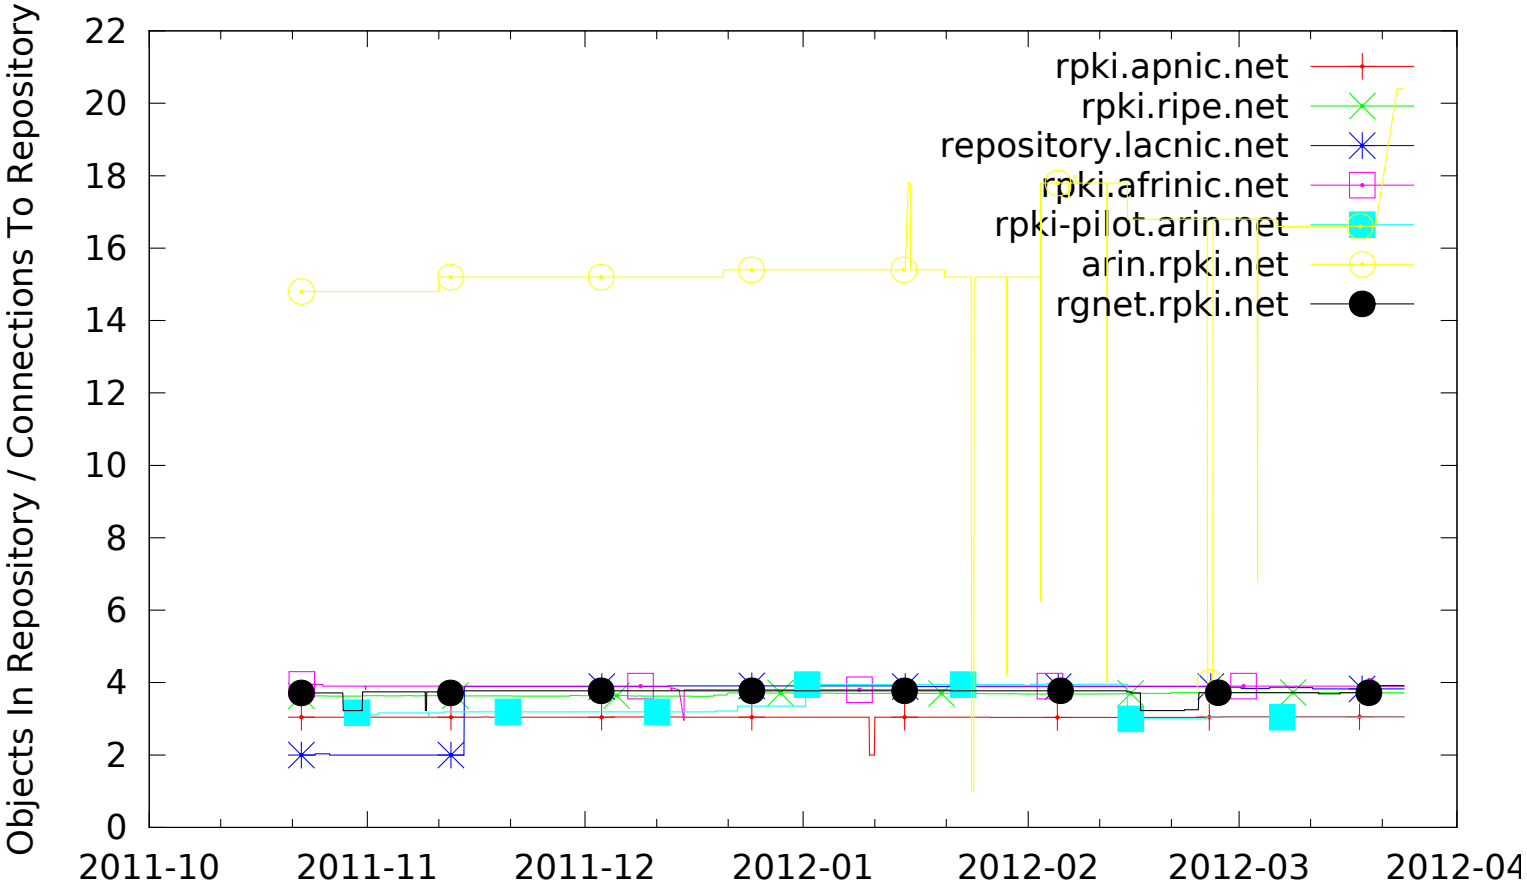

(Sessions with connection failures not shown)

Objects/Connection (Logarithmic)

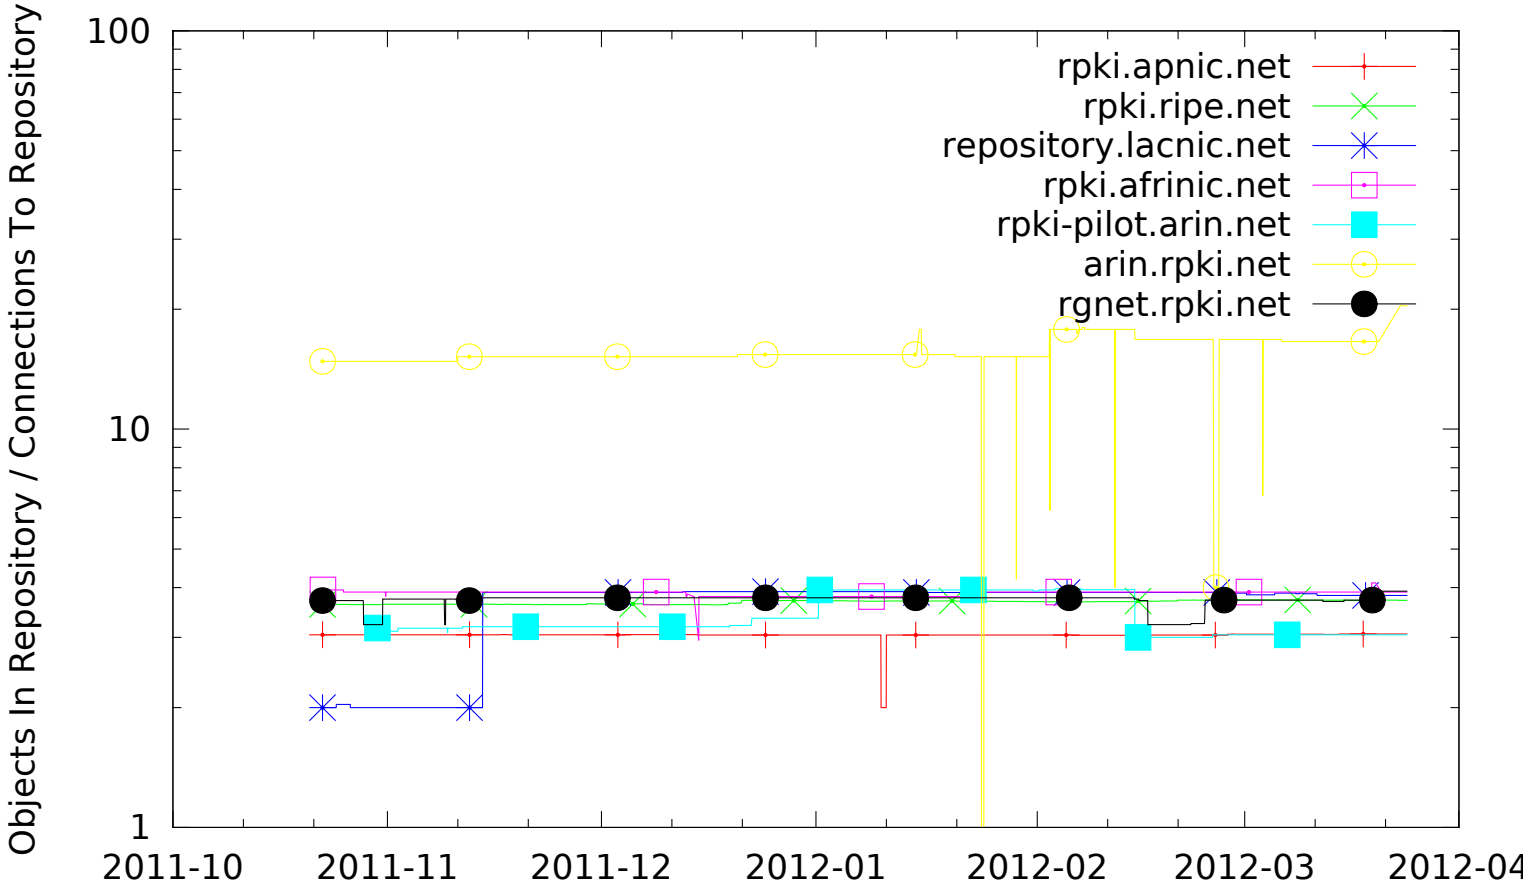

(Sessions with connection failures not shown)

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Seconds/Object (Linear)

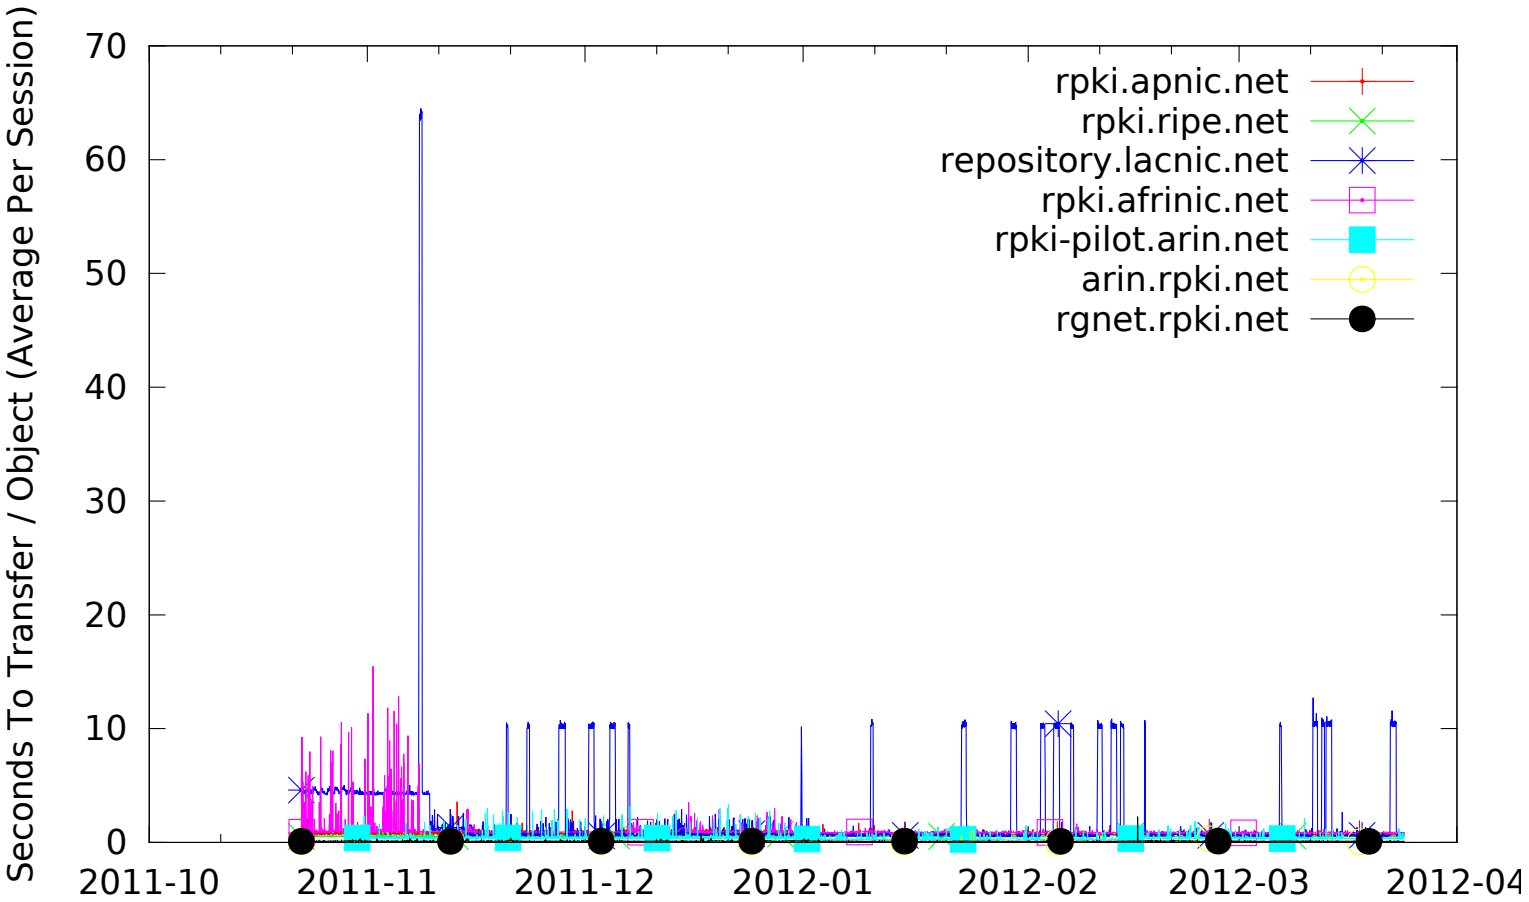

(Sessions with connection failures not shown)

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Seconds/Object (Logarithmic)

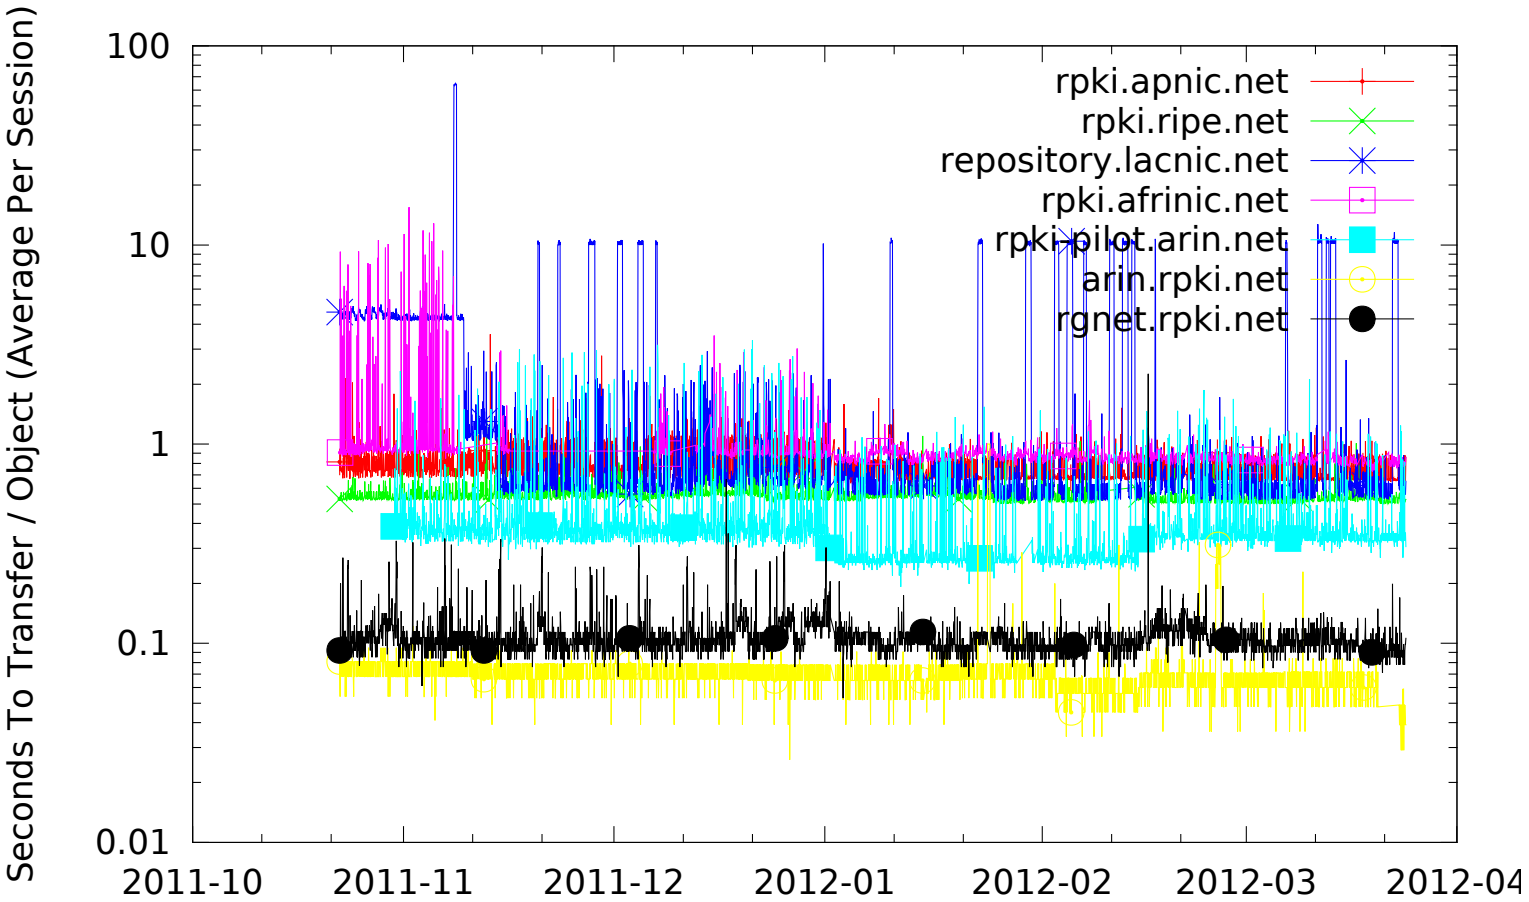

(Sessions with connection failures not shown)

Seconds/Object: Observations

- “Elapsed time” is sum of parallel connection times—five parallel connections of four minutes each counts as twenty minutes.
- We can speed up in terms of wall time by running more connections in parallel, but that puts more load on the repository servers and risks rate limiting (more on this later).
- Spikes here are slow repository servers; whether it’s the network path or the server itself that’s slow, we don’t know.

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Average Connection Duration (Linear)

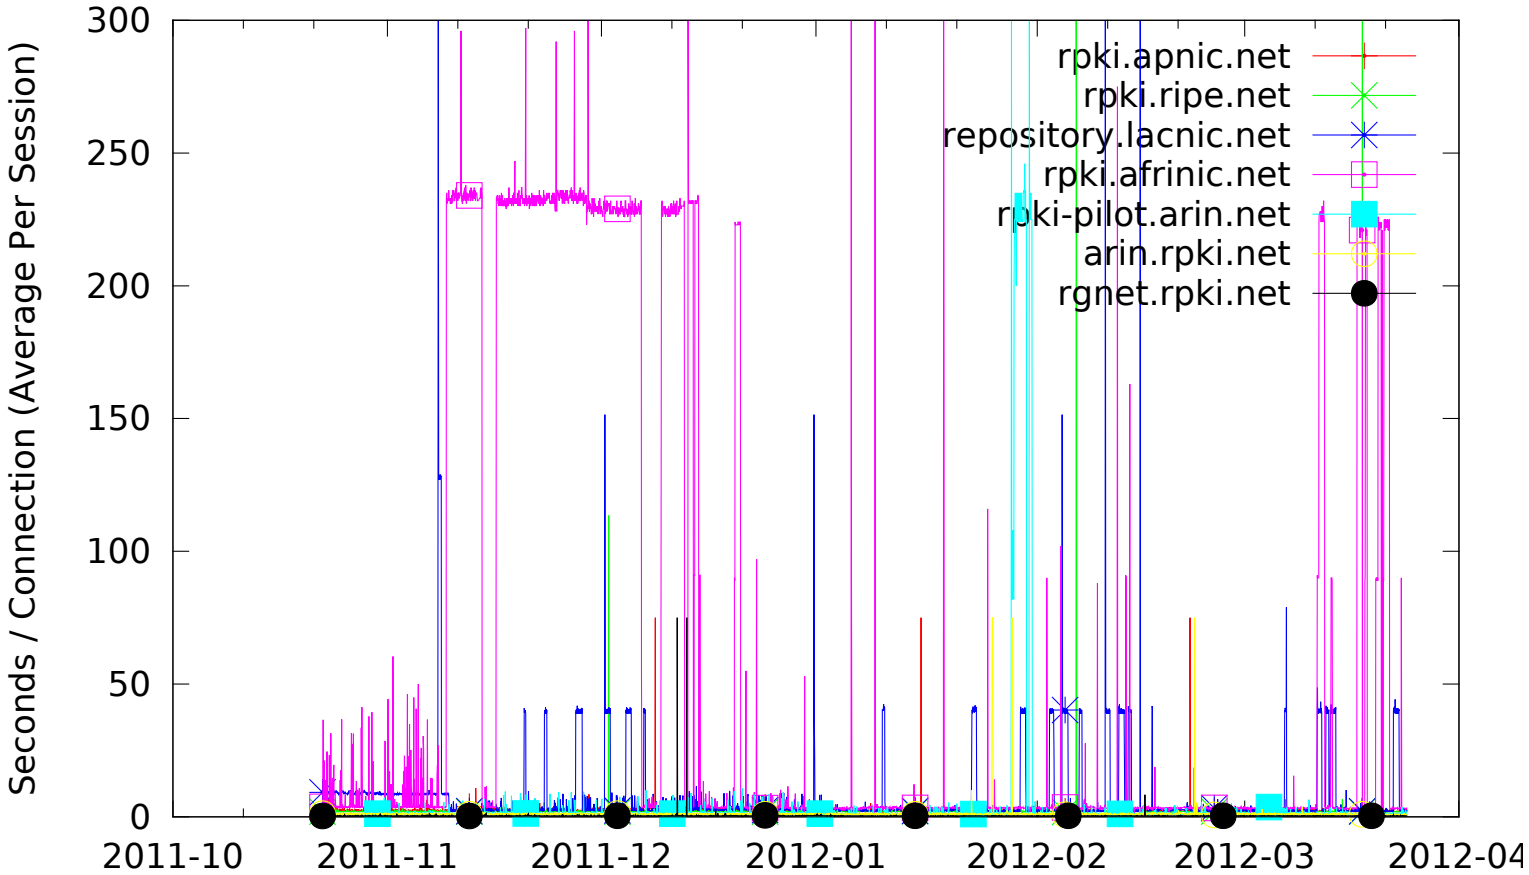

Average Connection Duration (Logarithmic)

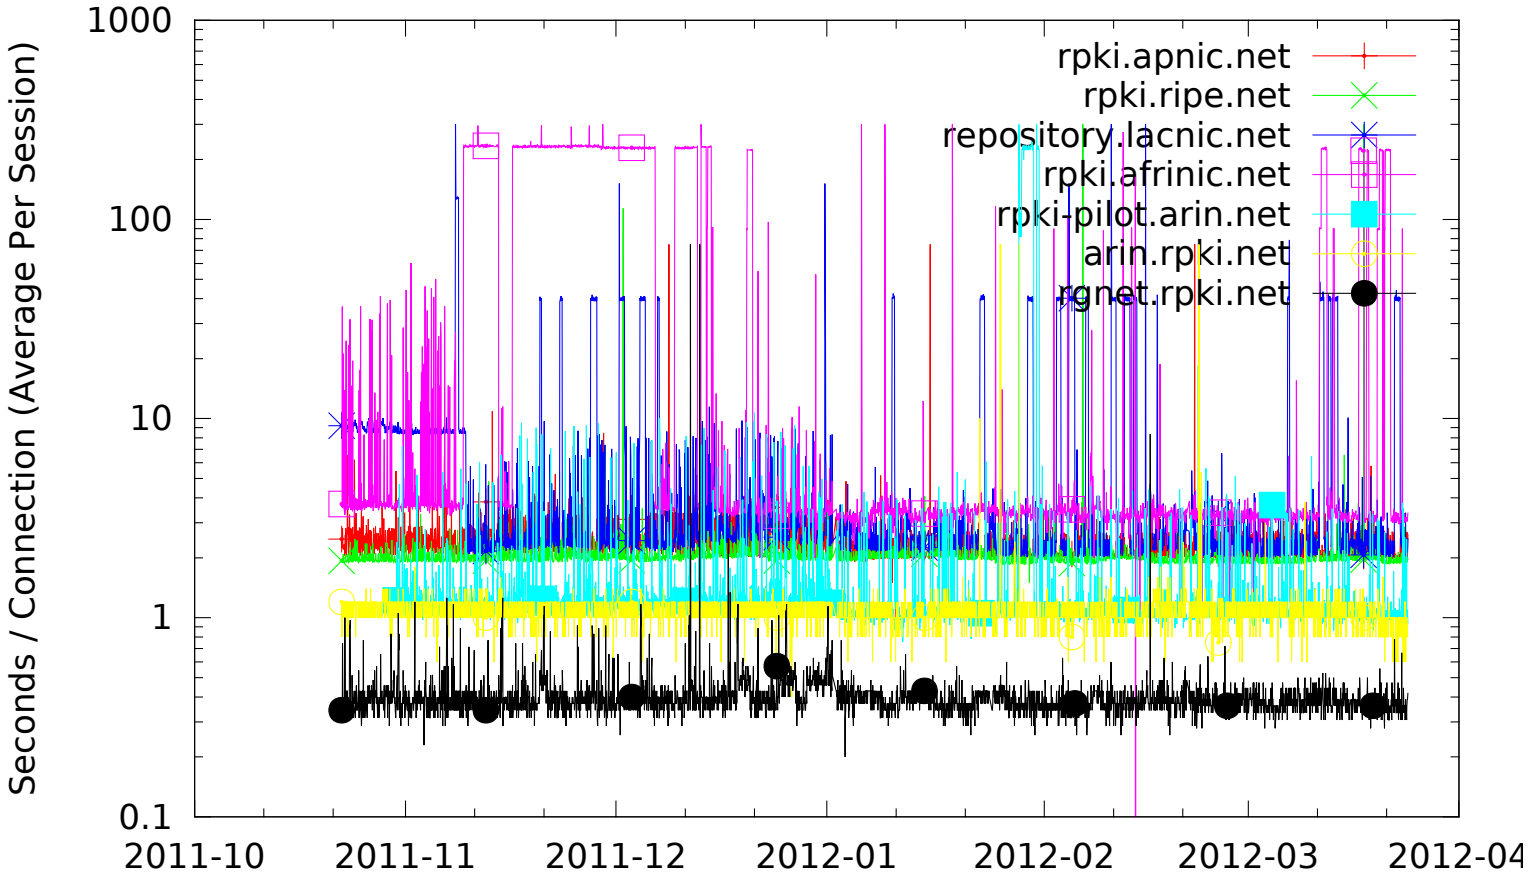

Average Connection Duration: Observations

- Early modeling and testing said much of cost is setup and teardown (about 500ms) and that this cost tends to dominate for large numbers of connections. So far, this analysis has held up pretty well.
- Spikes top out at 300 seconds because that's when rcynic gives up and whacks any rsync subprocess that appears to be completely stalled. This shouldn't happen, and generally indicates that repository server or network is badly messed up.

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Failure Rate (Linear)

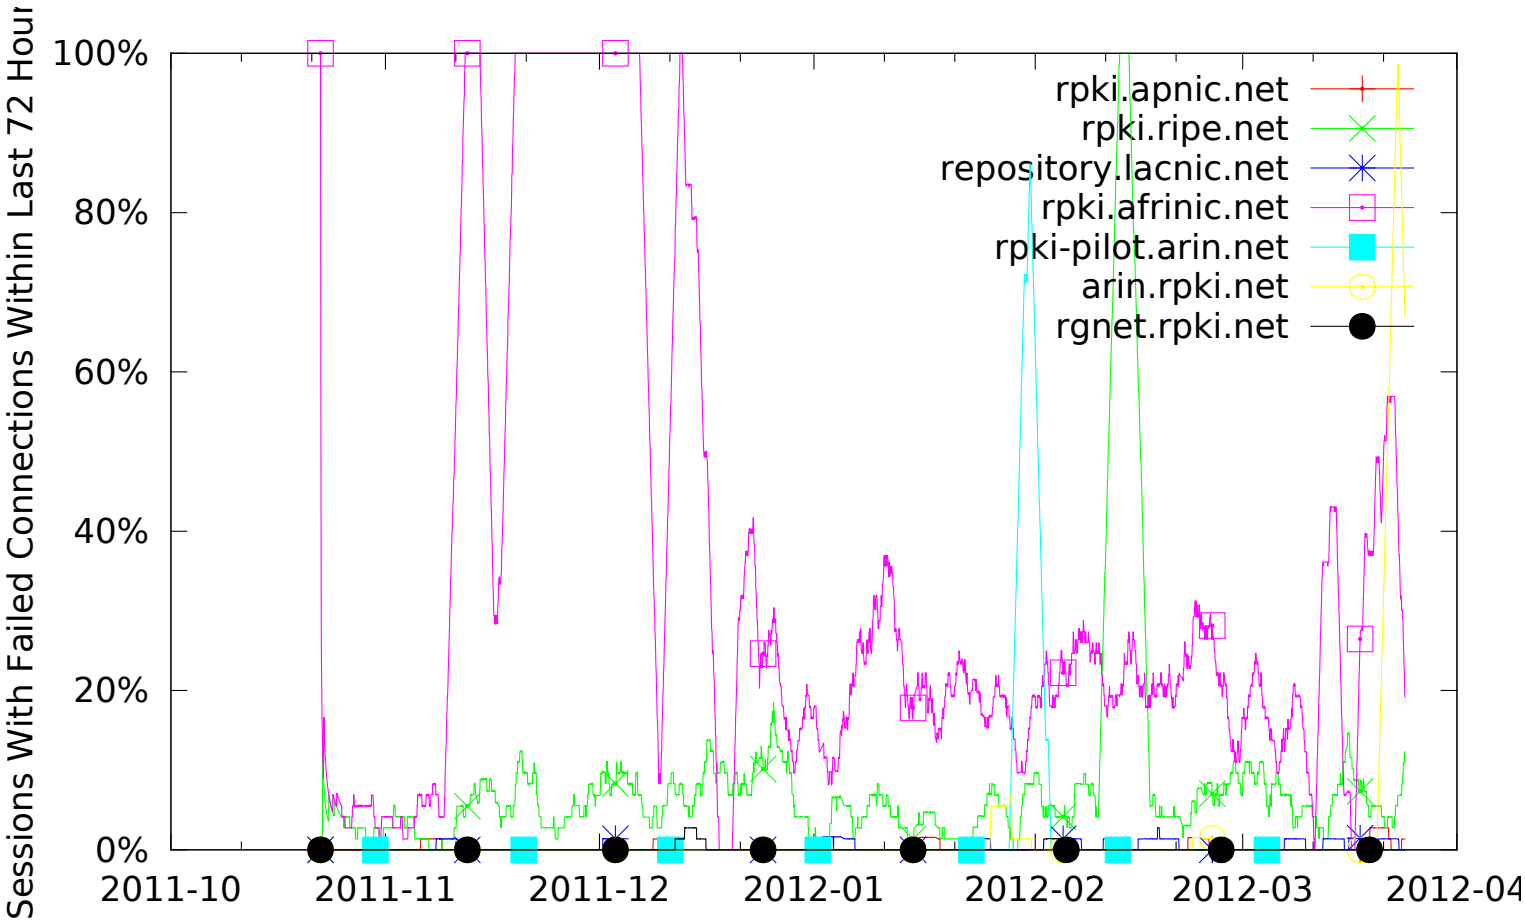

Failure Rate (Logarithmic)

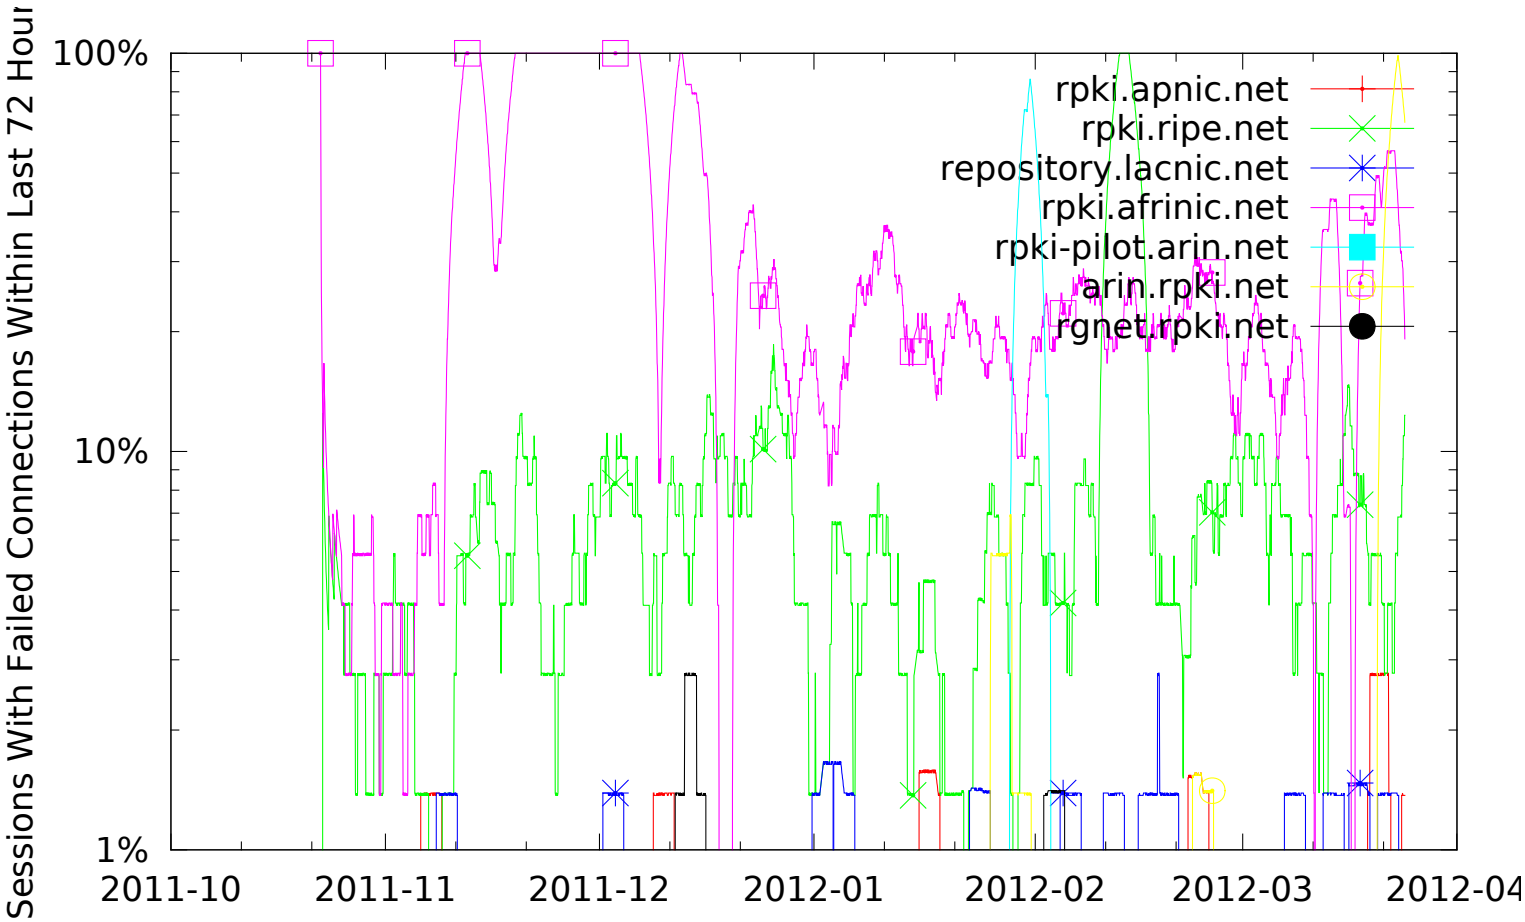

Failure Rate: Observations

- Failure rate is a bit hard to measure because:
  - We give up on a repository host for the duration of that session after the first failure.
  - rsync exit codes often don't tell us much we can use.
- For example, a valid certificate containing an incorrect SIA URI can result in a failure attempting to fetch from the named repository, with rsync exit code #23, "Partial transfer due to error."
- So shape of the curve is significant: a brief spike from 0% to 100% is probably a data error rather than a network issue, while a failure rate that wanders all over the map is probably a network or server.

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Rate Limiting (Sorry, No Graph)

- APNIC and AfriNIC used to rate limit to four connections in rsyncd.conf. Both appear to have stopped doing this.
- At one point APNIC also appeared to be rate limiting with some kind of firewall . . . which is harder to adapt to than rsyncd.conf limit. Haven't seen evidence of this recently.
- Others repositories currently appear to impose no rate limits.
- Rate limiting is a hard problem. What's the right limit for how many parallel rsync connections a validator should try? How should repository operator push back when overloaded?

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Sample Of Rcynic Status Output

- The following are samples of rcynic's normal output for the repositories in question.
- Some things are easier to see in this form, some are easier to see as graphs.
- We're still experimenting with how best to present these data.

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Summary for rpki.apnic.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																							442
current .cer														440	1							442	
current .crl														1								442	
current .mft														1								442	
current .roa												17										26	
Total												17		442	1							1352	442

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Summary for rpki.ripe.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																							925
current .cer														922	101							924	
current .crl														101								924	
current .mft														101	1							924	
backup .roa												2							2			2	
current .roa												659		175	17							666	
Total												661		1299	119				2			3440	925

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Summary for repository.lacnic.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																							57
current .cer														2	1							56	
current .crl														1								56	
current .mft														1								56	
current .roa																						50	
Total														4	1							218	57

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Summary for rpki.afrinic.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																				2			20
current .cer																					1	19	
current .crl																						19	
current .mft												1									1	19	
current .roa																						21	
Total												1								2	2	78	20

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Summary for rpki-pilot.arin.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																							44
current .cer															17		43		44			44	
current .crl																2						2	
current .mnf					44				44					17	14								
current .roa					44				44		4			14	16				44				
Total					88				88		4			31	47	2	43		88			46	44

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Summary for arin.rpki.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																							5
current .cer																	10	10	1			12	
backup .crl																1						1	
current .crl																7						8	
current .gbr				8					8		8						1	1	8			1	
current .mft																7	7					8	
backup .mnf																1	1					1	
backup .roa																	6	3	3			6	
current .roa		3		9		3		3	12		9						53	53	12			50	
Total		3		17		3		3	20		17					16	78	67	24			87	5

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Summary for rgnet.rpki.net 2012-03-24T22:04:50Z

	certificate has expired	certificate revoked	RFC 3779 resource not subset of parent's resources	AKI extension issuer mismatch	Bad keyUsage	Certificate failed validation	CRLDP doesn't match issuer's SIA	Manifest lists missing object	Object rejected	rsync transfer failed	AIA doesn't match issuer	EE certificate with 1024 bit key	Multiple rsync URIs in extension	Nonconformant X.509 issuer name	Nonconformant X.509 subject name	Stale CRL or manifest	Tainted by stale CRL	Tainted by stale manifest	Tainted by not being in manifest	Unknown object type skipped	Non-rsync URI in extension	Object accepted	rsync transfer succeeded
																							36
current .cer																						35	
current .crl																7						36	
current .gbr	1	1				2			2								1	1	3			3	
current .mft																7	7					36	
current .roa																	8	8	4			29	
Total	1	1				2			2							14	16	9	7			139	36

A Few Months In The Life Of An RPKI Validator

http://rpki.net/

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Things We're Not Measuring (Yet?)

Freshness: Some kind of measure of whether we're keeping up with what's being published, regardless of how we do it or how much pain is involved. One could make a case that this is the critical measurement and that all else is just dickering over the price.

What else?

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion

Problems We Think We're Seeing

- Slow repository servers are an issue for validator, whether they fail or not.
- Flat repository structure is an issue for validator.
- Rate limiting is an issue for validator and repository operator.
- Validator might not need to poll every URI every session.
- Alternate transports worth investigating (*e.g.* BitTorrent, separate presentation).

Introduction

Performance Graphs

Object Counts

Connection Counts

Objects/Connection

Seconds/Object

Average Connection Duration

Failure Rate

Rate Limiting

Repository Summaries

Conclusion

Questions?

Introduction

Performance Graphs

- Object Counts
- Connection Counts
- Objects/Connection
- Seconds/Object
- Average Connection Duration
- Failure Rate
- Rate Limiting

Repository Summaries

Conclusion