NIST
National Institute of
Standards and Technology

iTL INFORMATION TECHNOLOGY LABORATORY

Trustworthy Networking Program

RPKI
Authentication for BGP

Sebastian Spies

DENOG3 – 20.10.2011

NIST BGPSEC Project
O.Borchert, K.Lee, D.Montgomery, K.Sriram, O.Kim

Homeland Security
Science and Technology

Problem

- BGP Prefix Hijacking (for decades)
  - Youtube Incident
  - Table Leak of Chinanet (AS23734), ~37k routes
  - Pilosov/Kapela MITM Attack, many more

- BGP provides no way to
  - determine authorization of an AS to announce a prefix
  - validate path of a BGP update

What is Prefix Hijacking?

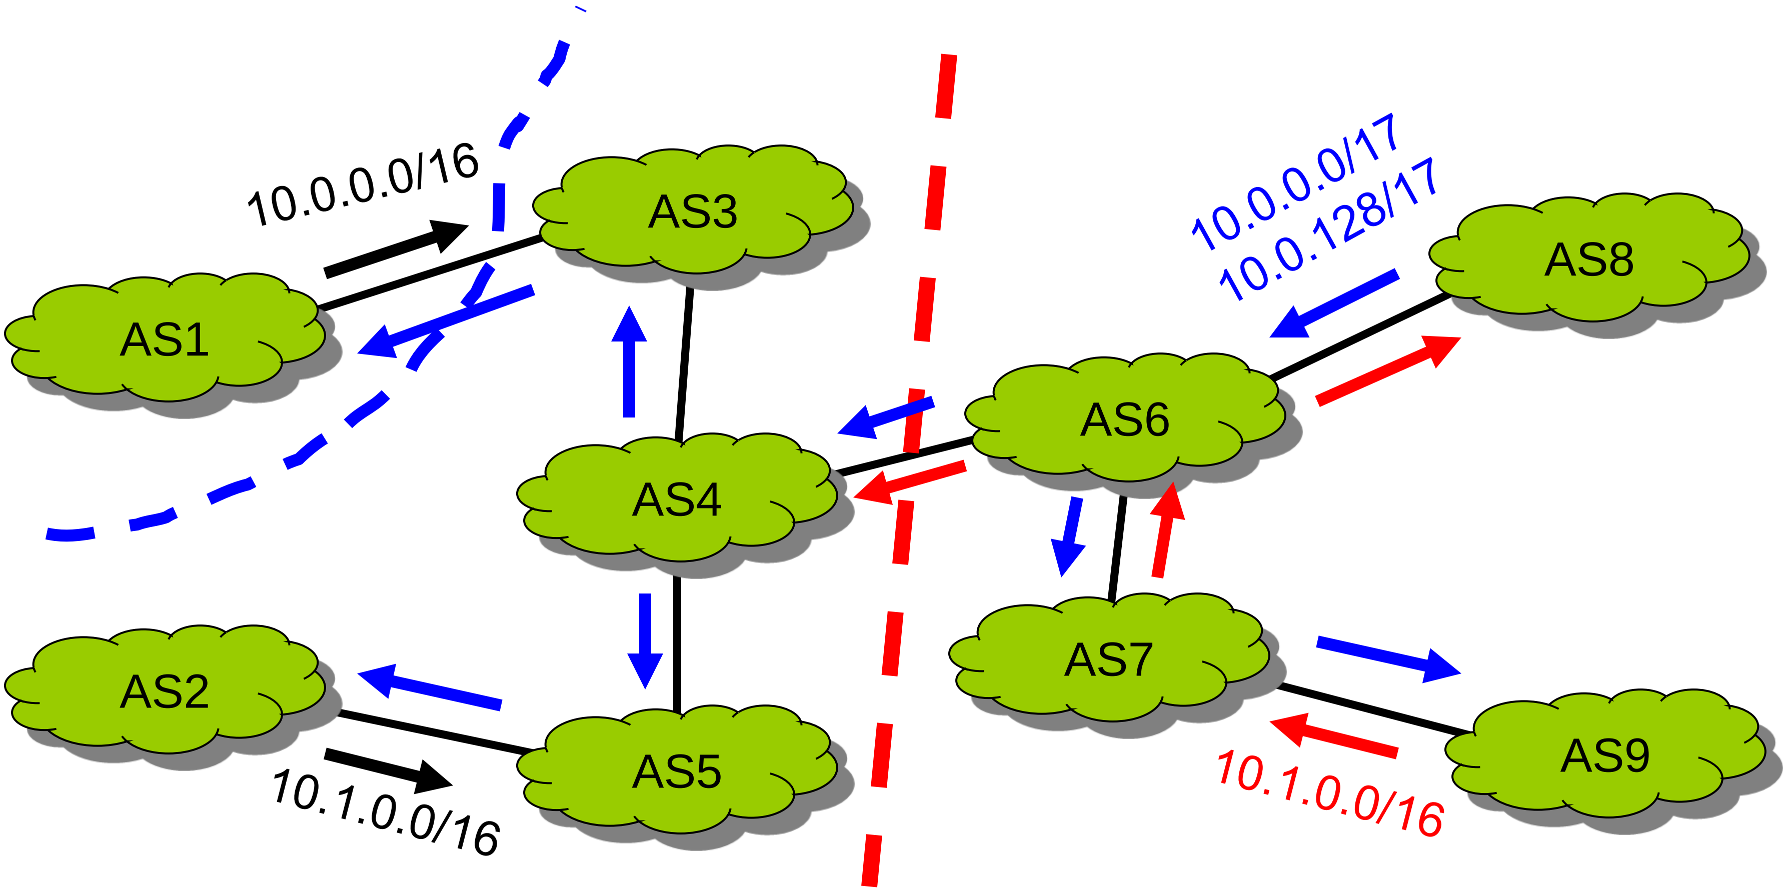

Invalid Announcement of prefix 10.1.0.0/16 cuts off AS6 – AS9

Invalid Announcement of more specific /17 prefix affects AS2-AS9

IETF SIDR WG
Proposed Solution

- Resource PKI (RPKI) enables routers to validate if the origin AS of a BGP update is correct (Route Origin Authorization, ROA)
- BGPSEC (with the help of RPKI) enables routers to cryptographically ensure, that a BGP update has traversed the ASNs in the path
- Take origin validation and BGPSEC to secure the control plane

What is the RPKI

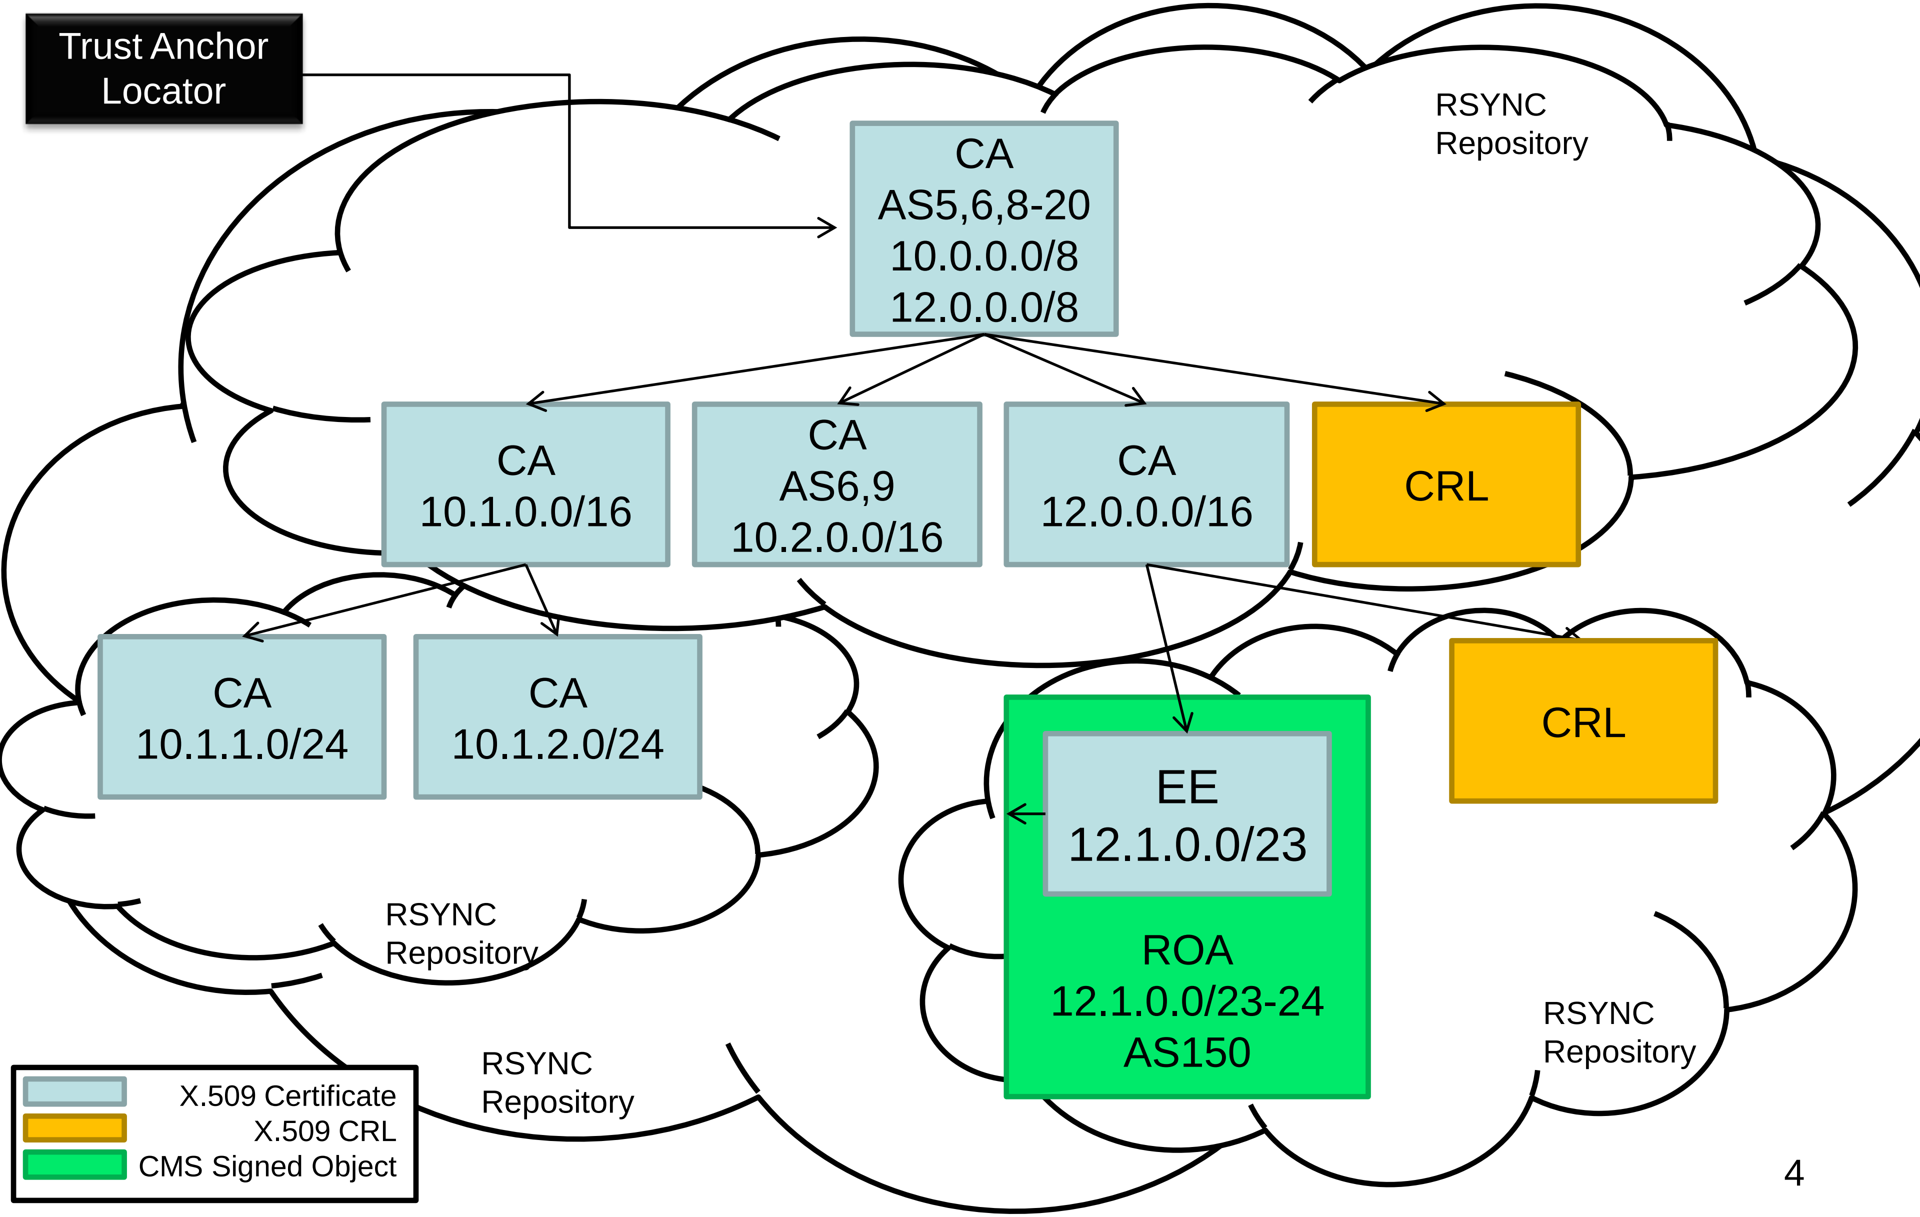

From the Repositories to the Router

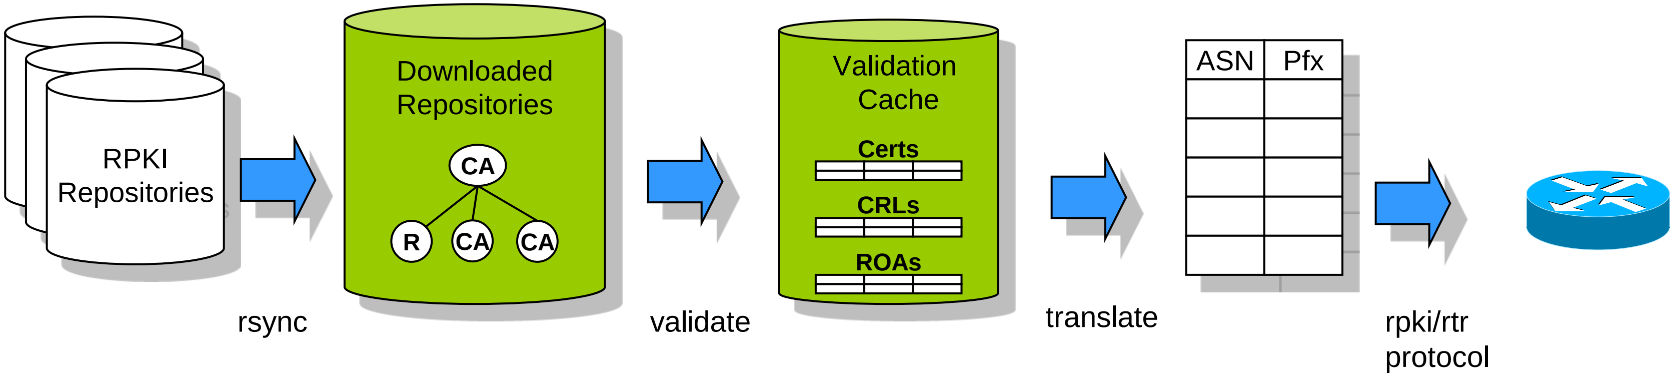

- Remote Synchronize the RPKI repositories into a local repository
- Validate ROAs with regards to expiry, sub-allocation, CRLs, etc.
- Translate valid ROAs into a prefix/origin list
- Communicate it to the router
- Validation tools from ISC, RIPE, BBN

Protocols

- Resource Cert Provisioning Protocol
  - aka “up/down“ protocol
  - Cert request, issuance, revocation and status info
  - HTTP POSTs of CMS signed-objects containing XML
  - Content-Type application/rpki-updown
- Publication Protocol
  - Like provisioning
  - For configuration of repository server and publish/withdraw certs to/from repository
  - Content-Type application/rpki-publication
- RPKI/RTR Protocol
  - Validation Cache sends prefix/originAS pairs to router
  - Incremental Updates
  - Transport Protocol
    - unprotected, TCP AO (preferred), SSH Transport Proto, TCP MD5, IPSec, TLS

RPKI/RTR Protocol

```
Cache                         Router
  ~                             ~
  | <----- Reset Query -------- | R requests data
  |                             |     (or Serial Query)
  | ----- Cache Response -----> | C confirms request
  | ------- IPvX Prefix ------> | C sends zero or more
  | ------- IPvX Prefix ------> | IPv4 and IPv6 Prefix
  | ------- IPvX Prefix ------> | Payload PDUs
  | ------ End of Data ------>  | C sends End of Data
  |                             | and sends new serial
  ~                             ~
  | -------- Notify ----------> |  (optional)
  |                             |
  | <----- Serial Query ------- | R requests data
  |                             |
  | ----- Cache Response -----> | C confirms request
  | ------- IPvX Prefix ------> | C sends zero or more
  | ------- IPvX Prefix ------> |   IPv4 and IPv6 Prefix
  | ------- IPvX Prefix ------> |   Payload PDUs
  | ------  End of Data ------> | C sends End of Data
  |                             |   and sends new serial
```

from draft-ietf-sidr-rpki-rtr-18

Origin Validation
States of a Route

- VALID
  ROA found, that matches routes‘ prefix and origin AS and satisfies maxlength
- INVALID
  There was at least one ROA, that matches prefix (regardless of maxlength), but none of them matches routes‘ origin AS and fits into maxlength
  (i.e. ROA 10.2.0.0/16-19 ASN5,
  Update 10.2.2.0/24 Origin AS5
  Update 10.2.2.0/17 Origin AS6)
- UNKNOWN/NOT FOUND
  There is no ROA, that matches the prefix of the route

BGPSEC Overview

- Assumes ROA and RPKI
- Cryptographic assurance of AS_PATH
- Router signs BGP updates
- Put AS number and router id into RPKI certs and deploy keys to routers
- (Unresolved) Issues
  - Optimization needed
  - Route Servers (transparent AS in path)
  - Proxy Aggregation (AS_SETs deprecated)
  - Rebeaconing (due to expiry time)
  - Only one prefix per update (NLRI unpacking)
  - Multiple Crypto Algorithms (RSA-2048, ECDSA-224, ECDSA-256)

BGPSEC Path Attribute Signature

```
Sequence of Octets to be Signed
when originating a route

+---------------------------------------+
| Expire Time (8 octets)                |
+---------------------------------------+
| Target AS Number (4 octets)           |
+---------------------------------------+
| Origin AS Number (4 octets)           |
+---------------------------------------+
| Algorithm Suite Identifier  (1 octet) |
+---------------------------------------+
| NLRI Length  (1 octet)                |
+---------------------------------------+
| NLRI Prefix  (variable)               |
+---------------------------------------+


Sequence of Octets to be Signed
when advertising a learned route

+-----------------------------------------------------------+
| Most Recent Signature Field   (fixed by algorithm suite)  |
------------------------------------------------------------+
| Target AS Number  (4 octets)                              |
+-----------------------------------------------------------+

from draft-ietf-sidr-bgpsec-protocol-00
```

BGPSEC RIB Size Estimation

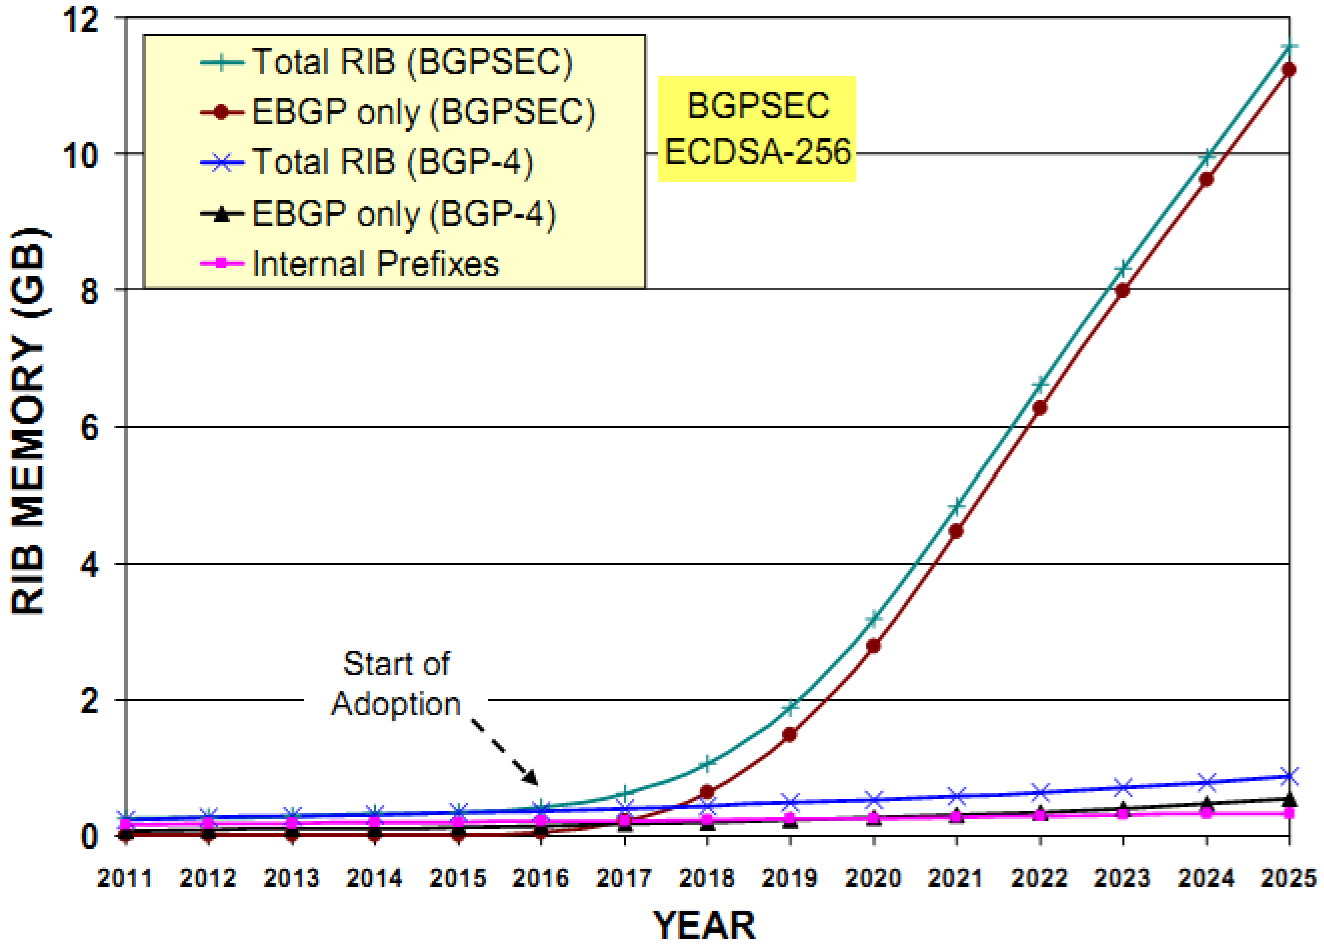

ECDSA-256

NIST Tools to Foster RPKI/BGPSEC Development

- BGP Secure Routing Extension (BGP-SRx)
  - Open Source Reference Implementation for RPKI processing within a router
  - Current stage – Prototype 0.2
    - BGP-SRx Server: Implementation talking to a validation cache using RPKI/RTR protocol
    - BGP-SRx API: Allows integration into BGP routers, policy modules, etc.
    - QuaggaSRx: Integrates BGP-SRx API into Quagga 0.99.16

- BGP RPKI Interoperability Tester and Evaluator (BRITE)
  - Web-based system, that tests
    - ROA Validation caches
    - BGP Routers, that use ROA Validation results using RPKI to router protocol

BGP-SRx Overview

- **Open Source Reference Implementation**
  - Software router with extensions for: RPKI cache maintenance, ROA and BGPSEC processing of updates, BGP route policies based upon new security tools.
  - BGP Secure Routing Extension (BGP-SRx) is designed as extension for Quagga routing platform. Designed to support other platforms (e.g., XORP, etc.)
  - Designed to support experimentation with different architectural configurations of SRx and RPKI components,
- **Status**
  - BGP-SRx framework with RPKI and ROA processing implemented.
  - Hooks for BGPSEC Path Validation ….

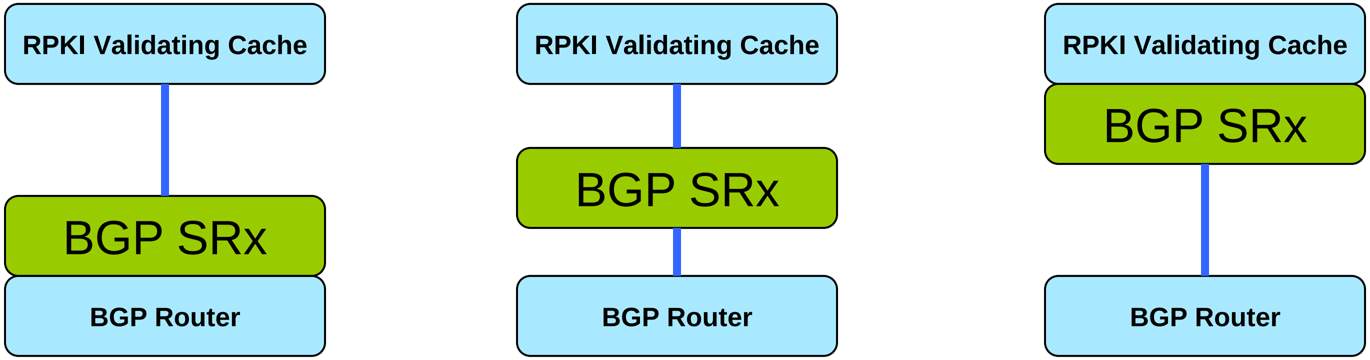

BGP-SRx System Architecture

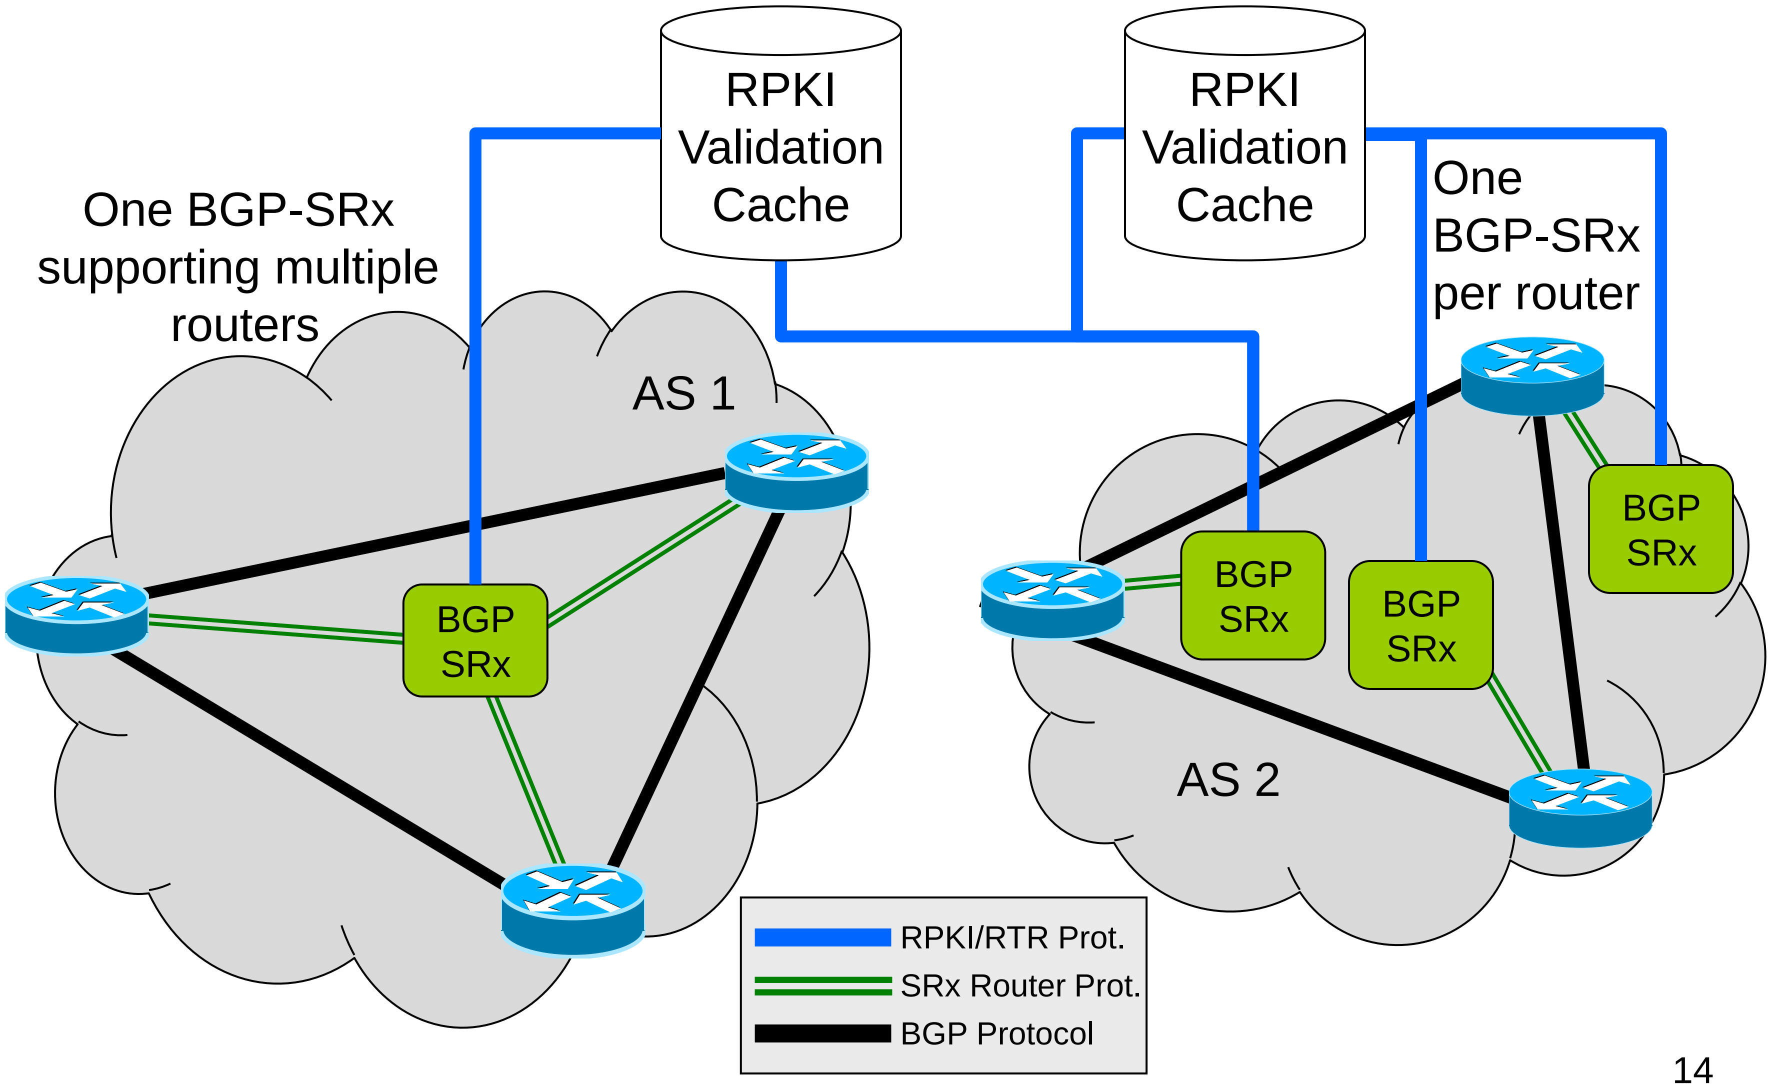

Quagga SRx Integration

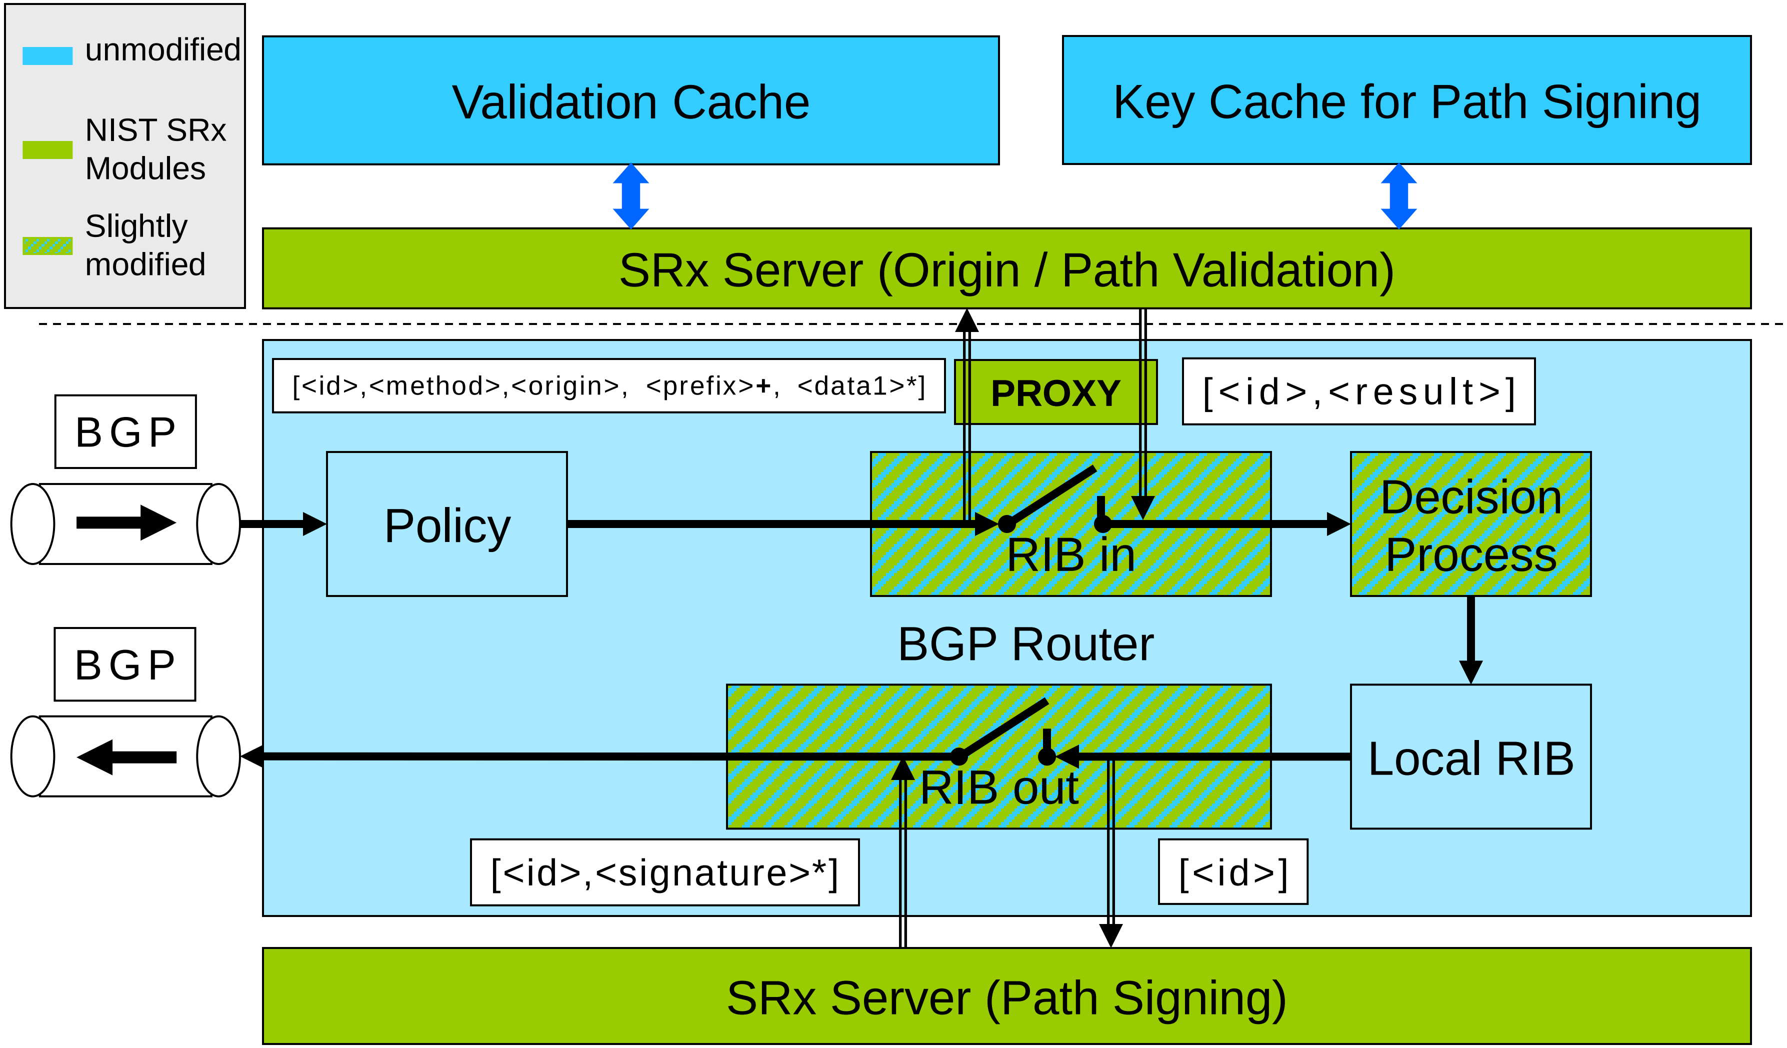

Quagga SRx Policy Set

- Activation of BGP-SRx Evaluation
  - no srx evaluation
  - srx evaluation (origin_only|bgpsec)
- Ignore Policies
  - [no] srx policy ignore-unknown
  - [no] srx policy ignore-invalid
  - [no] srx policy ignore-undefined
- Local Preference Policies
  - [no] srx policy local-preference valid <int> (add|subtract)
  - [no] srx policy local-preference unknown <int> (add|subtract)
  - [no] srx policy local-preference invalid <int> (add|subtract)
- Prefer Policies
  - [no] srx prefer-valid

BRITE Overview

- **BGPSEC / RPKI Interoperability Test & Evaluation**
  - Distributed test and evaluation framework for:
    - RPKI / BGP Security implementation testing,
    - Configuration and deployment testing.
  - Flexible XML based test / scenario scripting language.
  - Can test all components / interfaces of BGPSEC system.
    - RPKI Validating Caches
    - Cache to Router Protocol
    - ROA Processing in BGP Router

- **Distributed / automated test system.**
  - Webinterface to BRITE
  - Multi-user distributed architecture and interface
  - Real time test monitoring & reporting
  - Other diagnostics – log files, traffic traces available for download

Intention of BRITE

- BRITE is intended for
  - Developers of ROA validation/BGPSEC software as test bed
  - Early adopters to assess implications on their infrastructure
  - Operators
    - to verify test configuration settings
    - to be able to evaluate different RPKI/BGPSEC software packets
  - Researchers to study real-world behavior and stress test system configurations

BRITE Design Overview

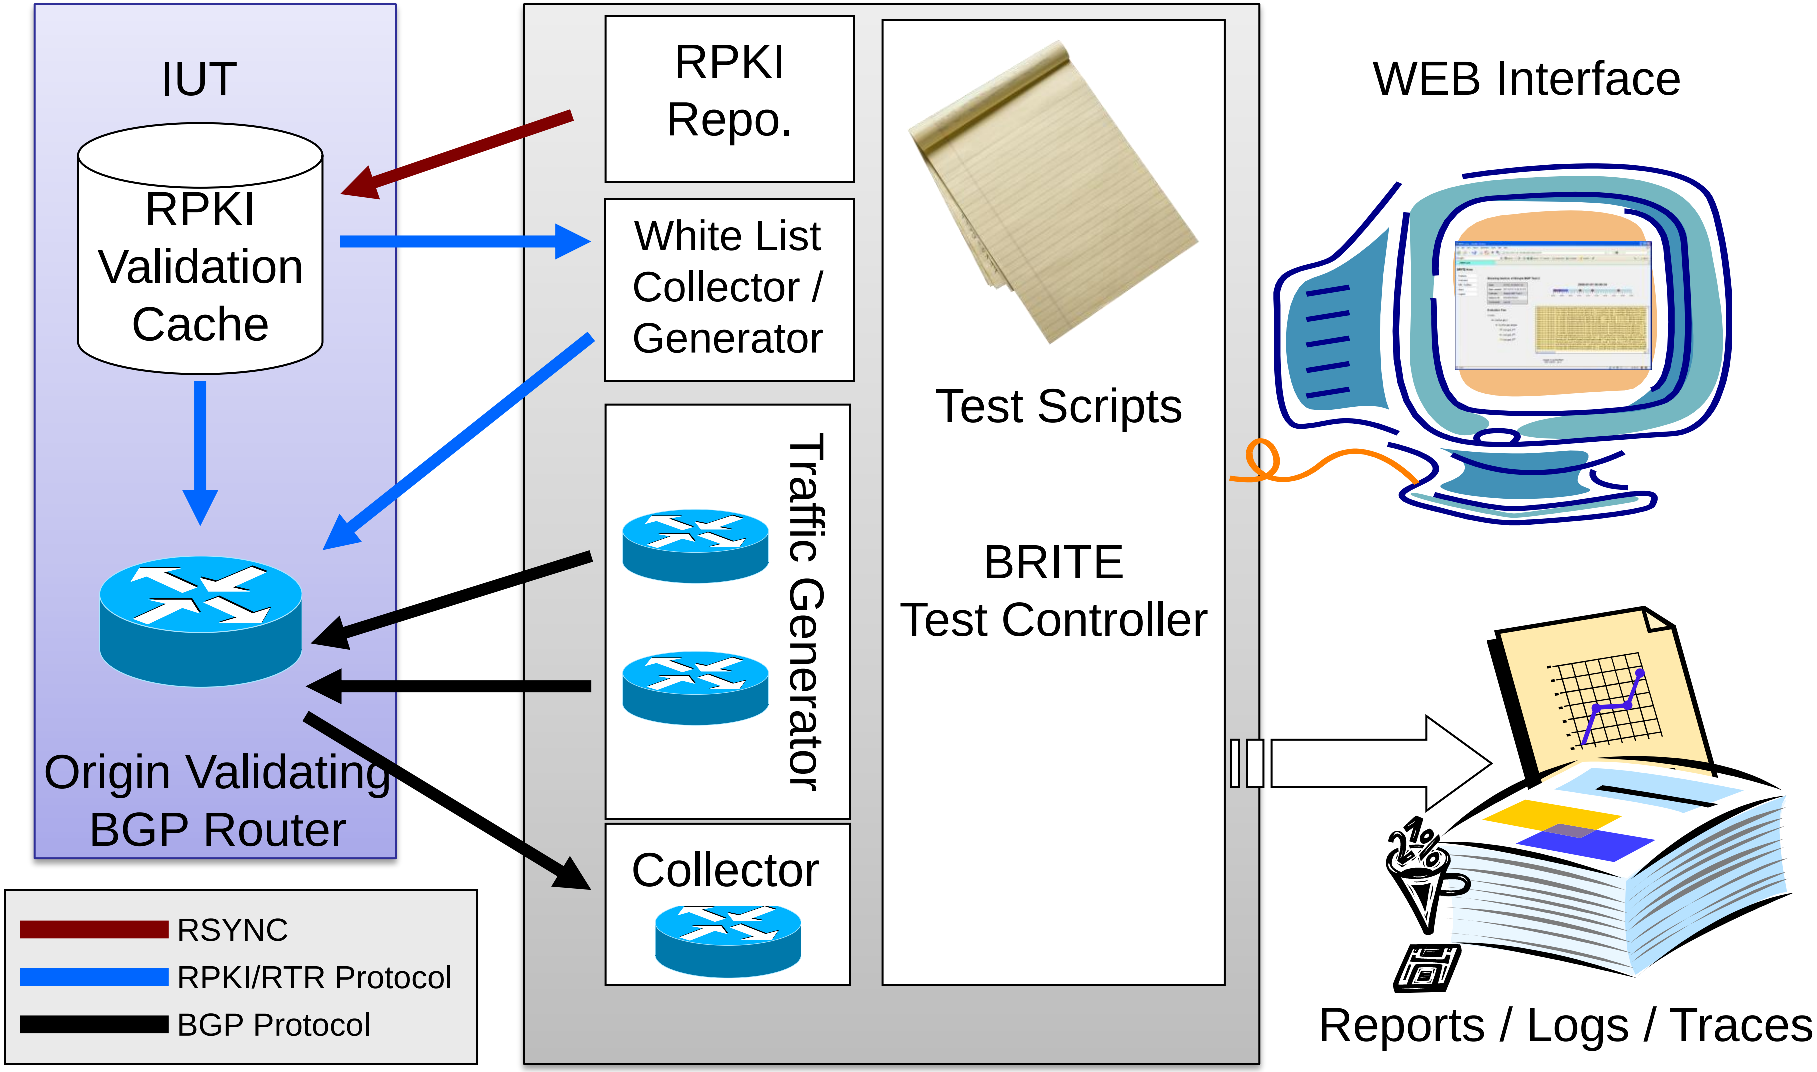

IUT

RPKI Validation Cache

Origin Validating BGP Router

RPKI Repo.

White List Collector / Generator

Traffic Generator

Collector

Test Scripts

BRITE Test Controller

WEB Interface

Reports / Logs / Traces

- RSYNC
- RPKI/RTR Protocol
- BGP Protocol

Demo – Simulated Topology

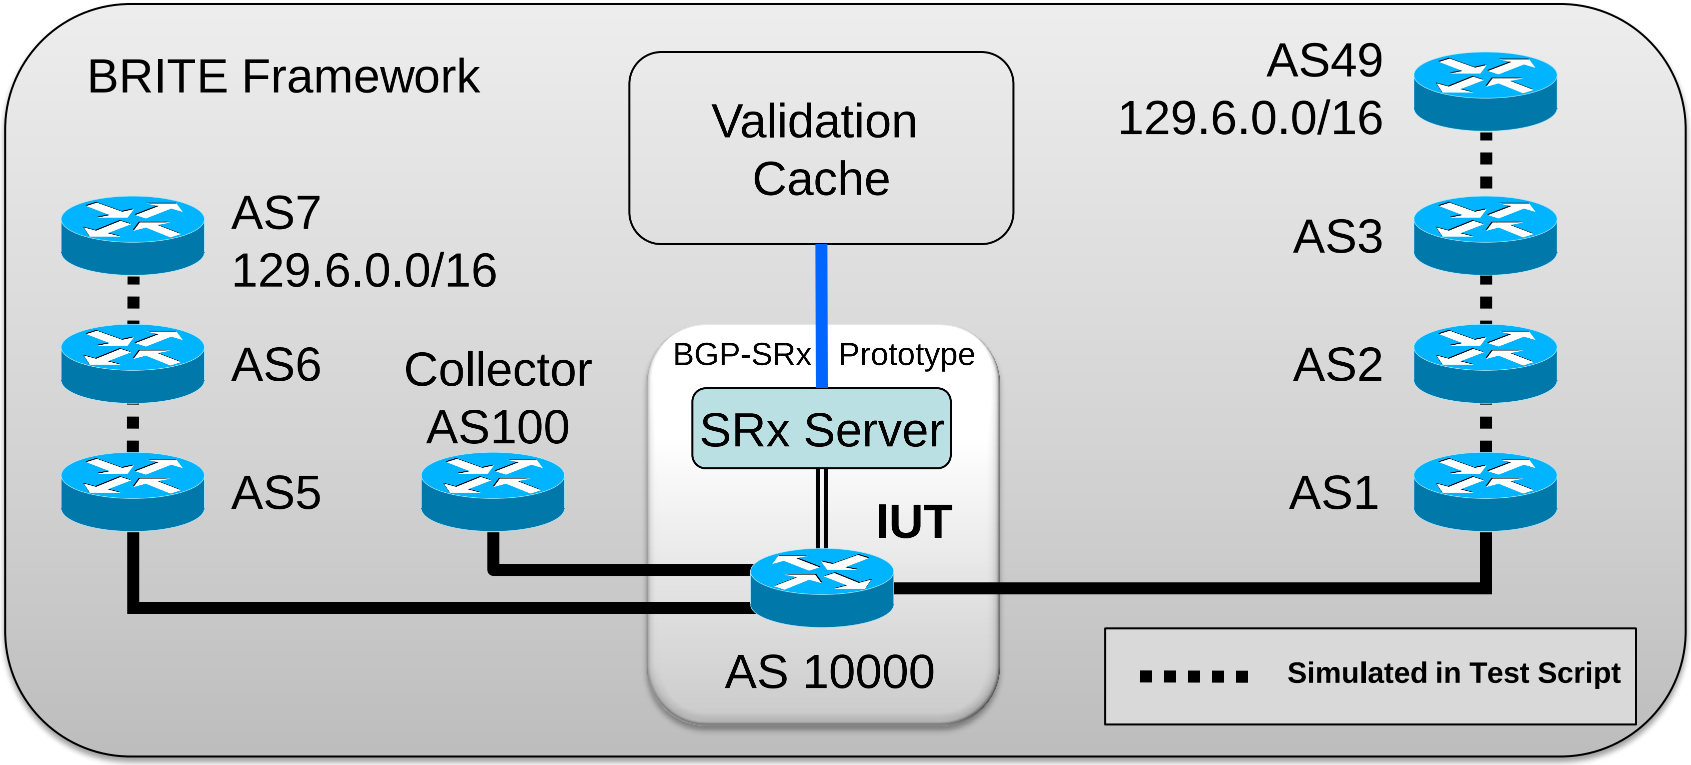

Test Event:

@t1: BGP: AS7 Originates 129.6.0.0/16
@t2: BGP: AS49 Originates129.6.0.0/16
@t3: RPKI: Add ROA {129.6.0.0/16-24, 49}
@t5: RPKI: Delete ROA {129.6.0.0/16-24, 49}

Test Goals (@collector):

@t1+: G1: BGP Ann. (129.6.0.0/16, AS7)

@t3+: G2: BGP Ann.(129.6.0.0/16, AS49)
@t5+: G3: BGP Ann.(129.6.0.0/16, AS7)

Thank you!

BGP – SRx

http://www-x.antd.nist.gov/bgpsrx

BRITE

http://brite.antd.nist.gov

Questions ?

XML Test Script Entities

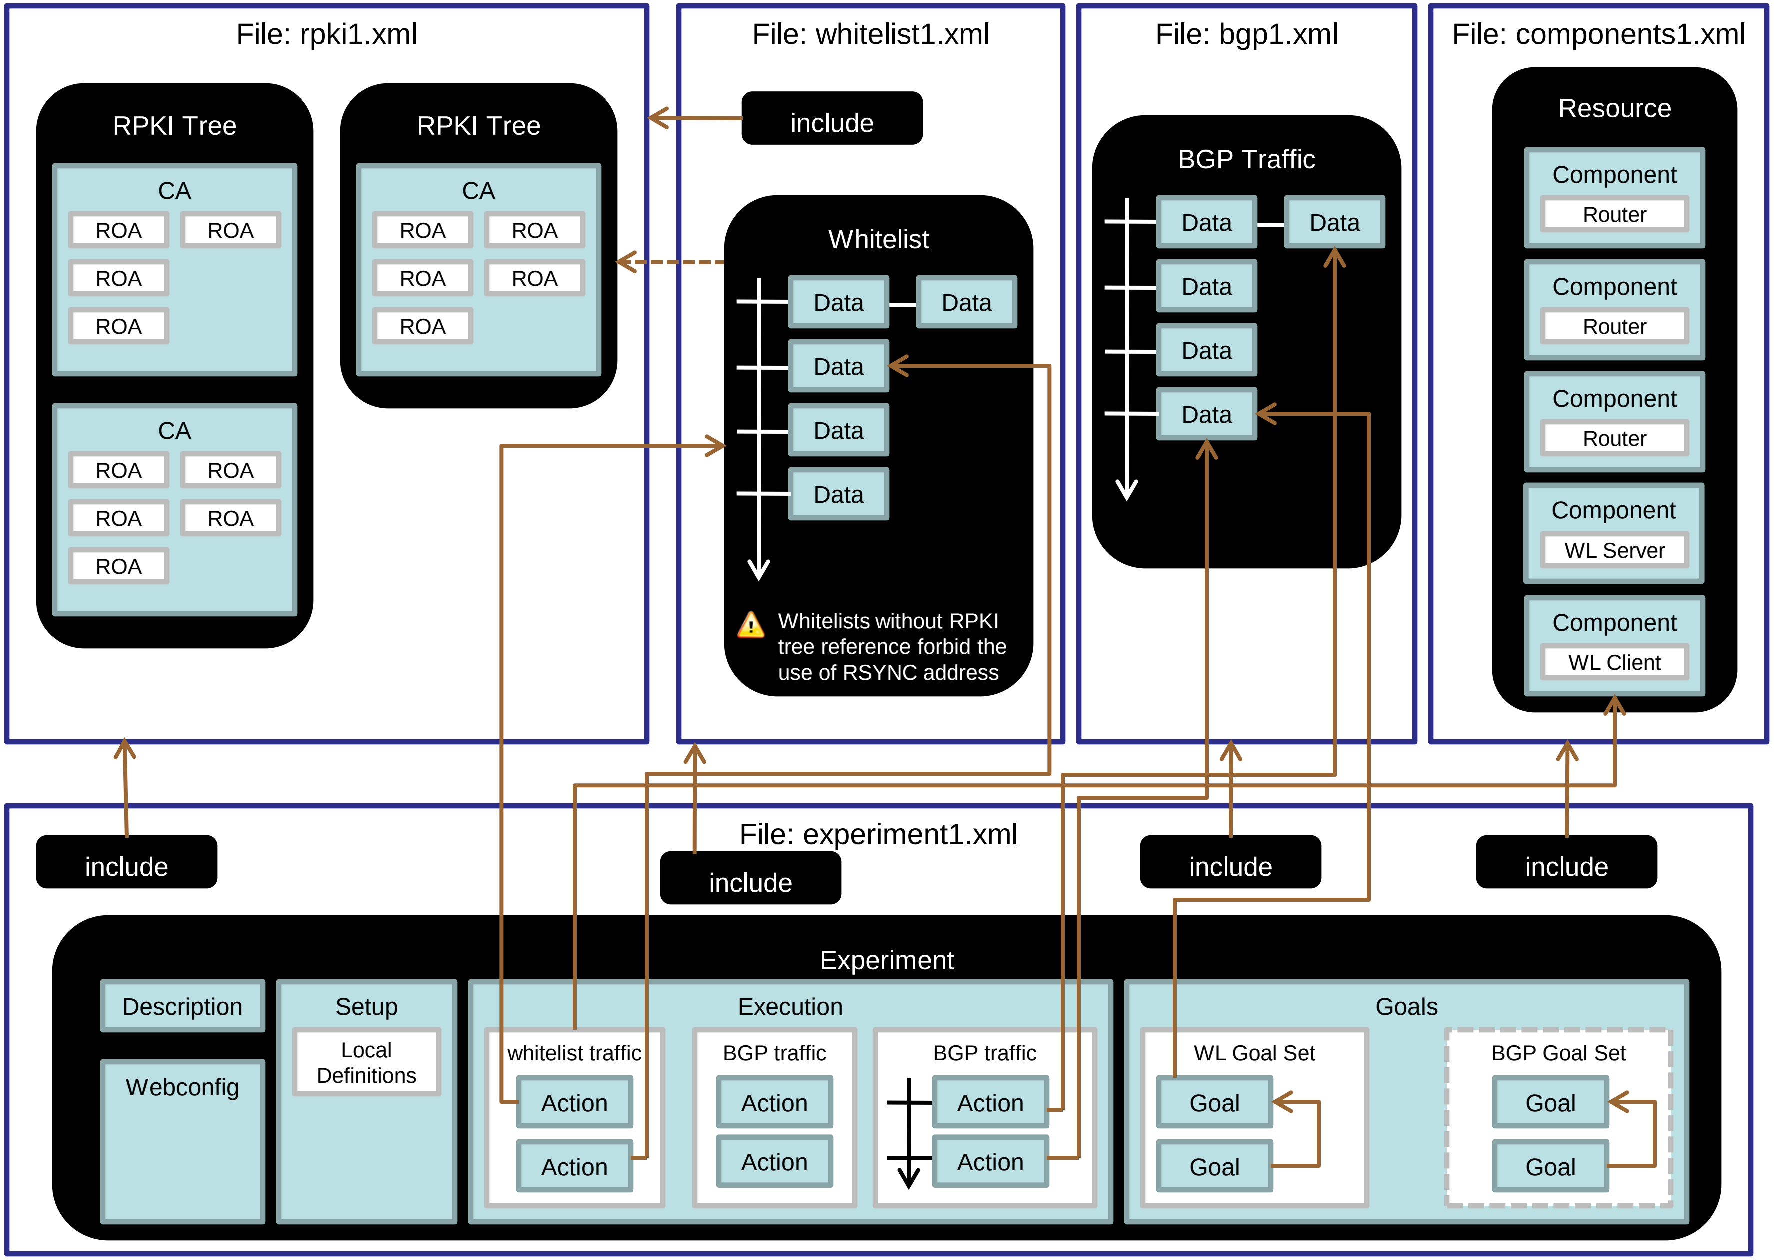

Remarks: Local definitions overwrite parents definition (local scope);

BRITE Web Interface (1)

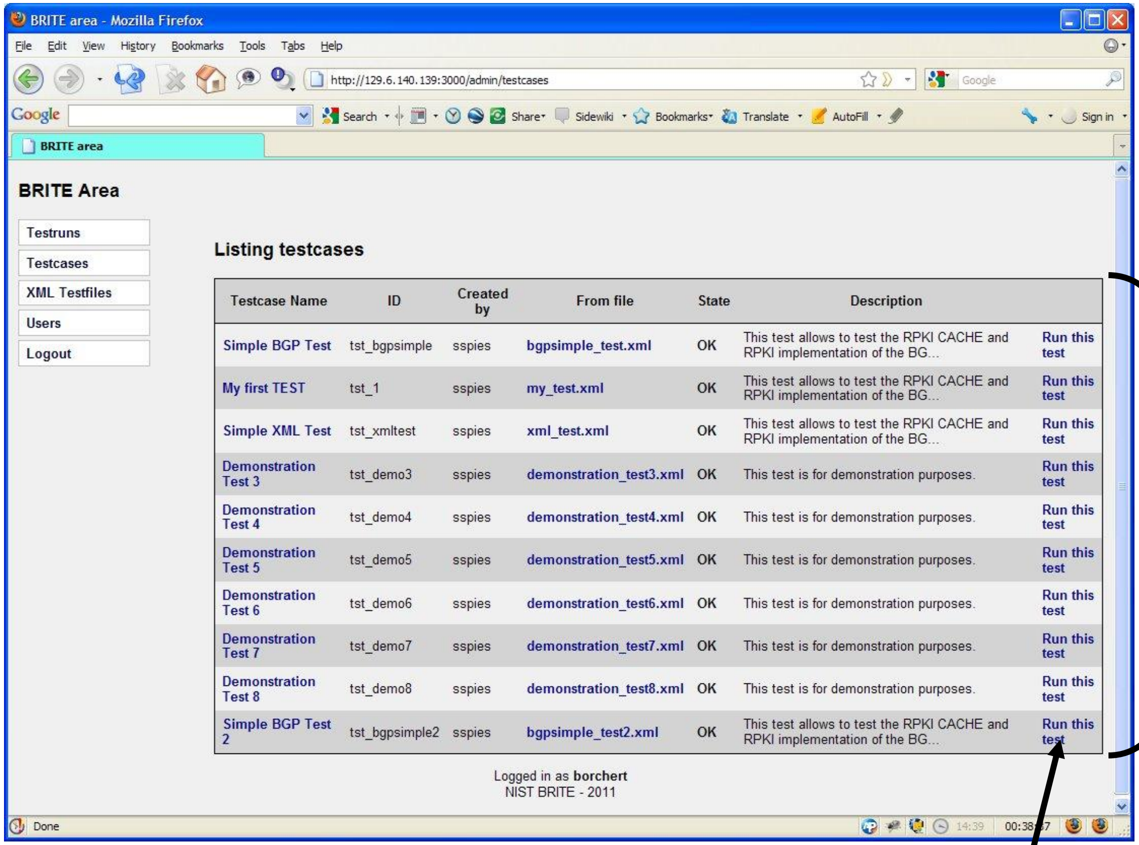
BRITE area - Mozilla Firefox
File Edit View History Bookmarks Tools Tabs Help
http://129.6.140.139:3000/admin/testcases
Google
Search Share Sidewiki Bookmarks Translate AutoFill Sign in
BRITE area
Done

BRITE Area

- Testruns
- Testcases
- XML Testfiles
- Users
- Logout

Listing testcases

Testcase Name	ID	Created by	From file	State	Description	
Simple BGP Test	tst_bgpsimple	sspies	bgpsimple_test.xml	OK	This test allows to test the RPKI CACHE and RPKI implementation of the BG...	Run this test
My first TEST	tst_1	sspies	my_test.xml	OK	This test allows to test the RPKI CACHE and RPKI implementation of the BG...	Run this test
Simple XML Test	tst_xmltest	sspies	xml_test.xml	OK	This test allows to test the RPKI CACHE and RPKI implementation of the BG...	Run this test
Demonstration Test 3	tst_demo3	sspies	demonstration_test3.xml	OK	This test is for demonstration purposes.	Run this test
Demonstration Test 4	tst_demo4	sspies	demonstration_test4.xml	OK	This test is for demonstration purposes.	Run this test
Demonstration Test 5	tst_demo5	sspies	demonstration_test5.xml	OK	This test is for demonstration purposes.	Run this test
Demonstration Test 6	tst_demo6	sspies	demonstration_test6.xml	OK	This test is for demonstration purposes.	Run this test
Demonstration Test 7	tst_demo7	sspies	demonstration_test7.xml	OK	This test is for demonstration purposes.	Run this test
Demonstration Test 8	tst_demo8	sspies	demonstration_test8.xml	OK	This test is for demonstration purposes.	Run this test
Simple BGP Test 2	tst_bgpsimple2	sspies	bgpsimple_test2.xml	OK	This test allows to test the RPKI CACHE and RPKI implementation of the BG...	Run this test

Logged in as **borchert**
NIST BRITE - 2011

Tests available to the User.

Select a test to be started

BRITE Web Interface (2)

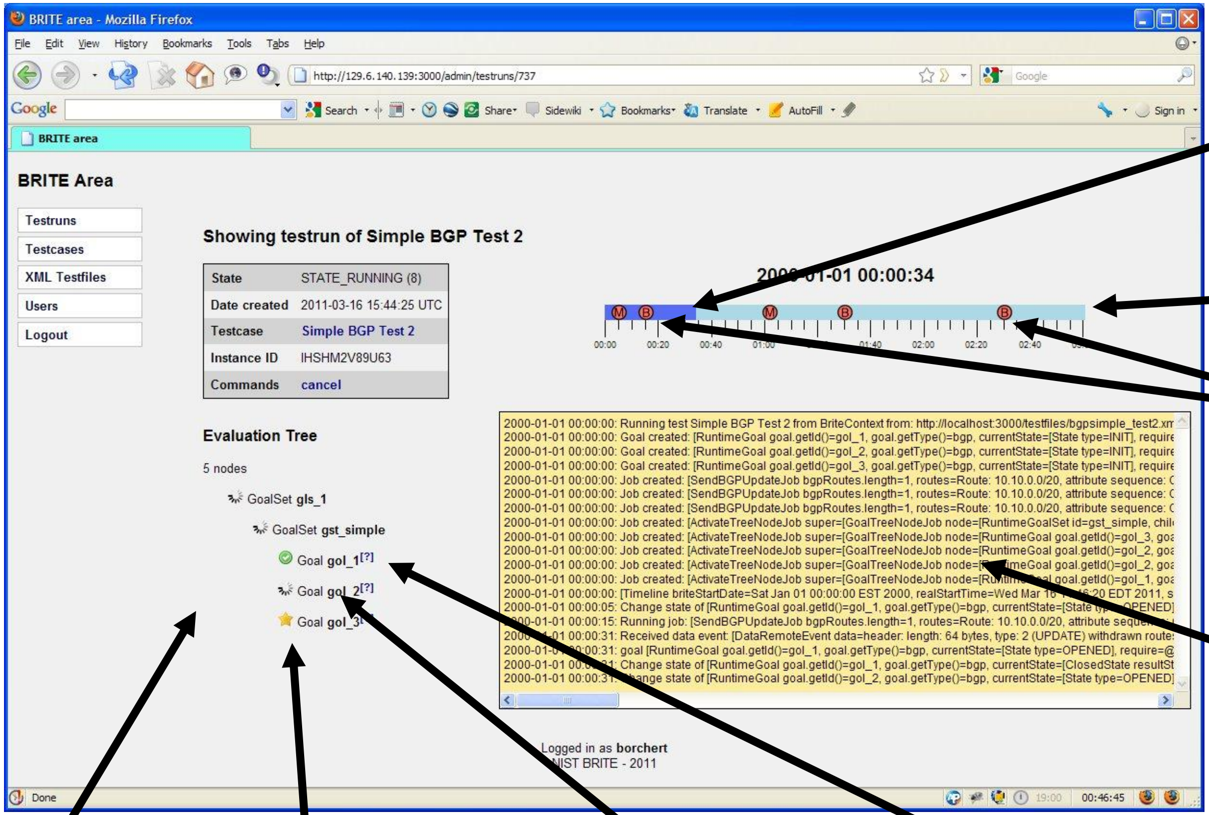

BRITE Area

Showing testrun of Simple BGP Test 2

State	STATE_RUNNING (8)
Date created	2011-03-16 15:44:25 UTC
Testcase	Simple BGP Test 2
Instance ID	IHSHM2V89U63
Commands	cancel

Evaluation Tree

5 nodes

- GoalSet **gls_1**
  - GoalSet **gst_simple**
    - Goal **gol_1**[?]
    - Goal **gol_2**[?]
    - Goal **gol_3**[?]

```
2000-01-01 00:00:00: Running test Simple BGP Test 2 from BriteContext from: http://localhost:3000/testfiles/bgpsimple_test2.xm
2000-01-01 00:00:00: Goal created: [RuntimeGoal goal.getId()=gol_1, goal.getType()=bgp, currentState=[State type=INIT], require
2000-01-01 00:00:00: Goal created: [RuntimeGoal goal.getId()=gol_2, goal.getType()=bgp, currentState=[State type=INIT], require
2000-01-01 00:00:00: Goal created: [RuntimeGoal goal.getId()=gol_3, goal.getType()=bgp, currentState=[State type=INIT], require
2000-01-01 00:00:00: Job created: [SendBGPUpdateJob bgpRoutes.length=1, routes=Route: 10.10.0.0/20, attribute sequence: (
2000-01-01 00:00:00: Job created: [SendBGPUpdateJob bgpRoutes.length=1, routes=Route: 10.10.0.0/20, attribute sequence: (
2000-01-01 00:00:00: Job created: [SendBGPUpdateJob bgpRoutes.length=1, routes=Route: 10.10.0.0/20, attribute sequence: (
2000-01-01 00:00:00: Job created: [ActivateTreeNodeJob super=[GoalTreeNodeJob node=[RuntimeGoalSet id=gst_simple, chil
2000-01-01 00:00:00: Job created: [ActivateTreeNodeJob super=[GoalTreeNodeJob node=[RuntimeGoal goal.getId()=gol_3, goa
2000-01-01 00:00:00: Job created: [ActivateTreeNodeJob super=[GoalTreeNodeJob node=[RuntimeGoal goal.getId()=gol_2, goa
2000-01-01 00:00:00: Job created: [ActivateTreeNodeJob super=[GoalTreeNodeJob node=[RuntimeGoal goal.getId()=gol_2, goa
2000-01-01 00:00:00: Job created: [ActivateTreeNodeJob super=[GoalTreeNodeJob node=[RuntimeGoal goal.getId()=gol_1, goa
2000-01-01 00:00:00: [Timeline briteStartDate=Sat Jan 01 00:00:00 EST 2000, realStartTime=Wed Mar 16 15:46:20 EDT 2011, s
2000-01-01 00:00:05: Change state of [RuntimeGoal goal.getId()=gol_1, goal.getType()=bgp, currentState=[State type=OPENED]
2000-01-01 00:00:15: Running job: [SendBGPUpdateJob bgpRoutes.length=1, routes=Route: 10.10.0.0/20, attribute seque
2000-01-01 00:00:31: Received data event: [DataRemoteEvent data=header: length: 64 bytes, type: 2 (UPDATE) withdrawn route:
2000-01-01 00:00:31: goal [RuntimeGoal goal.getId()=gol_1, goal.getType()=bgp, currentState=[State type=OPENED], require=@
2000-01-01 00:00:31: Change state of [RuntimeGoal goal.getId()=gol_1, goal.getType()=bgp, currentState=[ClosedState resultSt
2000-01-01 00:00:31: Change state of [RuntimeGoal goal.getId()=gol_2, goal.getType()=bgp, currentState=[State type=OPENED]
```

Logged in as **borchert**
NIST BRITE - 2011

Test Progress

Test Timeline

Events:
M=Multiple
A =Activation
B =BGP
W=Whitelist

Experiment Log

Goal Tree

Wait to be activated

Currently processing

Finished successful

QuaggaSRx (1)

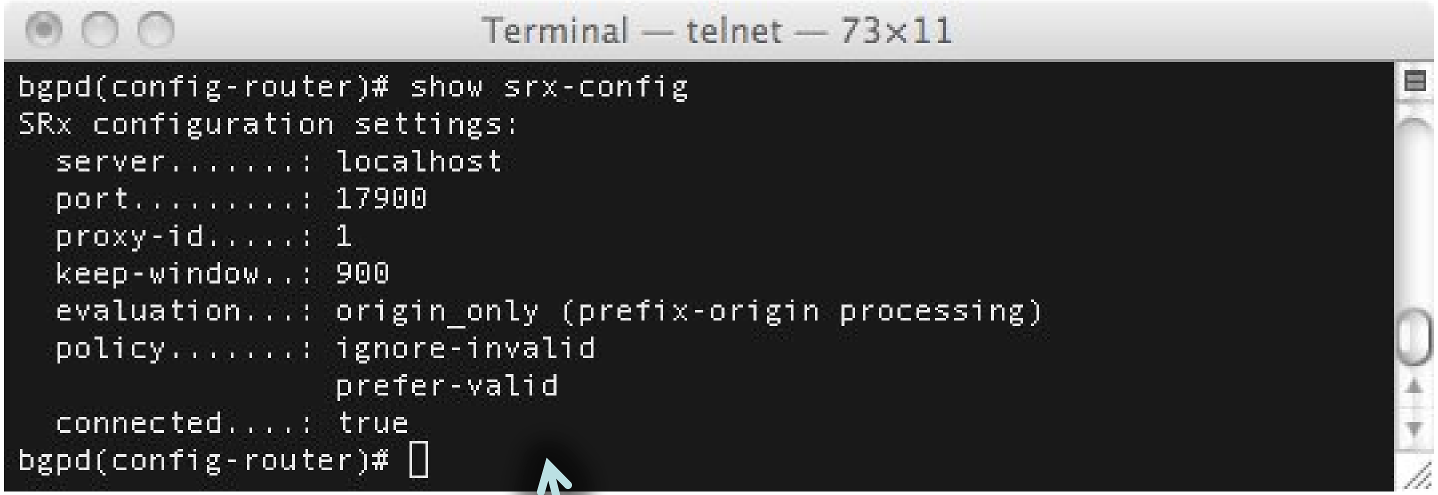

```
bgpd(config-router)# show srx-config
SRx configuration settings:
  server.......: localhost
  port.........: 17900
  proxy-id.....: 1
  keep-window..: 900
  evaluation...: origin_only (prefix-origin processing)
  policy.......: ignore-invalid
                 prefer-valid
  connected....: true
bgpd(config-router)# 
```

Configuration information related to SRx integration and origin / path processing!

QuaggaSRx (2)

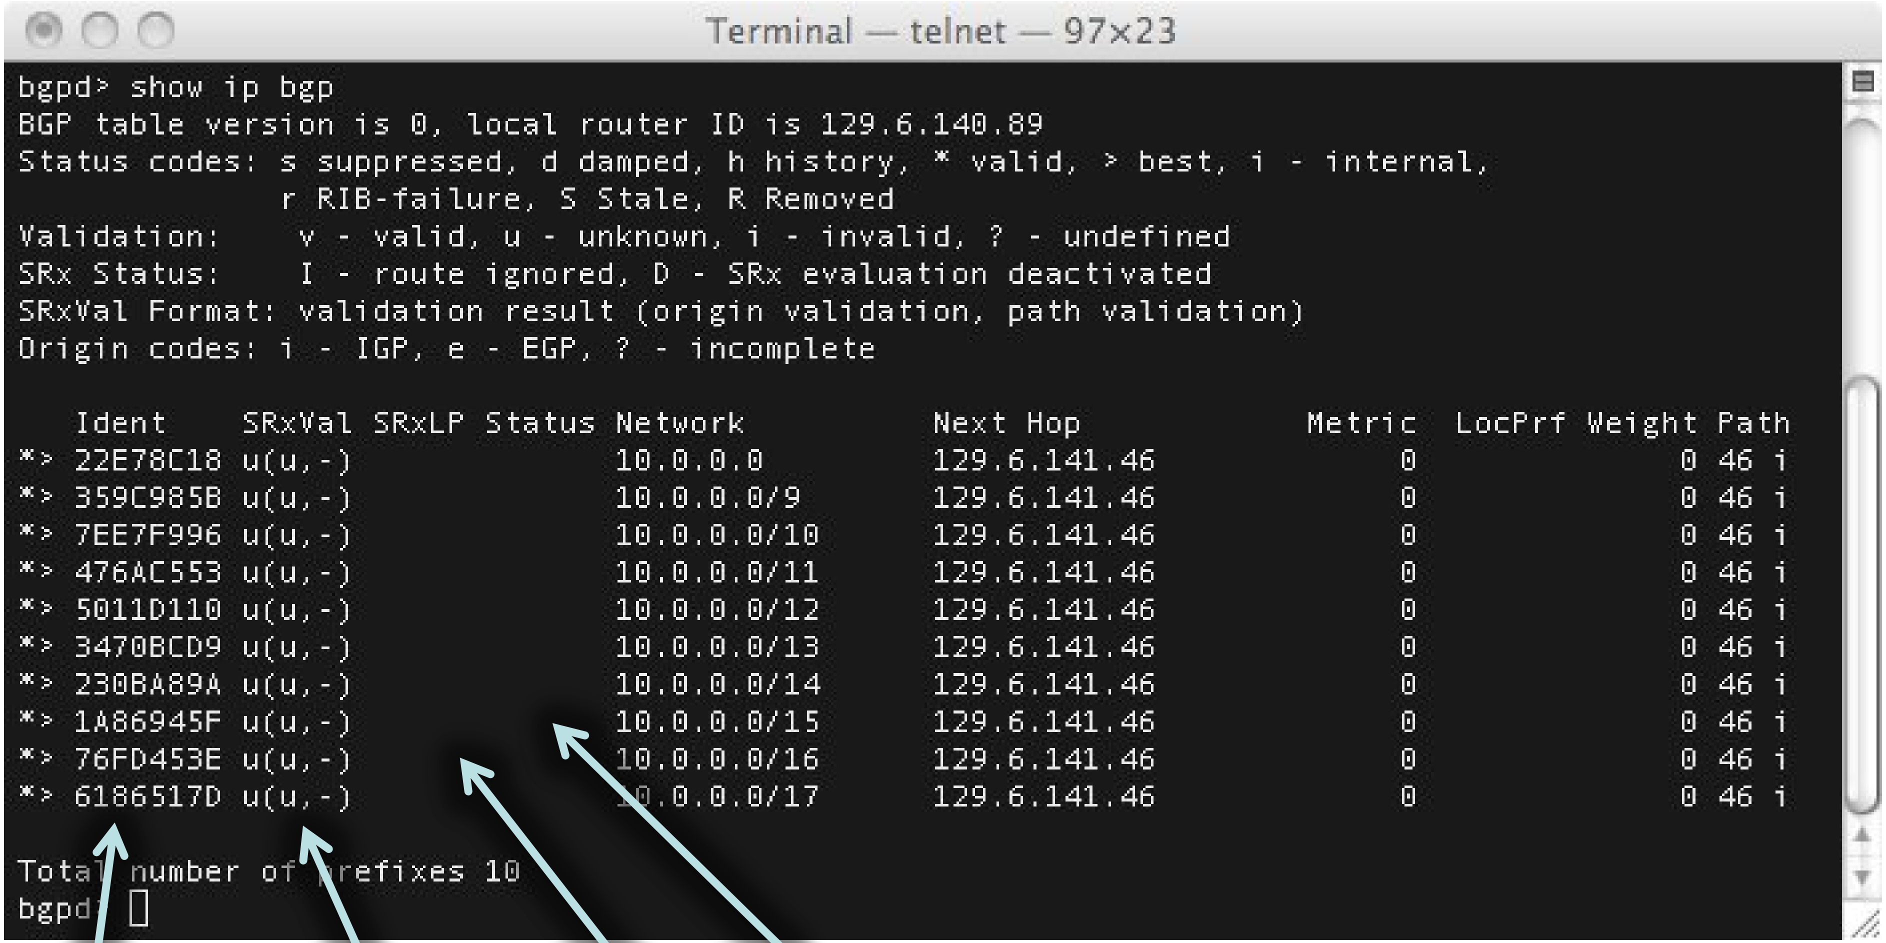

```
bgpd> show ip bgp
BGP table version is 0, local router ID is 129.6.140.89
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, R Removed
Validation:     v - valid, u - unknown, i - invalid, ? - undefined
SRx Status:     I - route ignored, D - SRx evaluation deactivated
SRxVal Format: validation result (origin validation, path validation)
Origin codes: i - IGP, e - EGP, ? - incomplete

   Ident    SRxVal SRxLP Status Network          Next Hop            Metric  LocPrf Weight Path
*> 22E78C18 u(u,-)              10.0.0.0         129.6.141.46             0              0 46 i
*> 359C985B u(u,-)              10.0.0.0/9       129.6.141.46             0              0 46 i
*> 7EE7F996 u(u,-)              10.0.0.0/10      129.6.141.46             0              0 46 i
*> 476AC553 u(u,-)              10.0.0.0/11      129.6.141.46             0              0 46 i
*> 5011D110 u(u,-)              10.0.0.0/12      129.6.141.46             0              0 46 i
*> 3470BCD9 u(u,-)              10.0.0.0/13      129.6.141.46             0              0 46 i
*> 230BA89A u(u,-)              10.0.0.0/14      129.6.141.46             0              0 46 i
*> 1A86945F u(u,-)              10.0.0.0/15      129.6.141.46             0              0 46 i
*> 76FD453E u(u,-)              10.0.0.0/16      129.6.141.46             0              0 46 i
*> 6186517D u(u,-)              10.0.0.0/17      129.6.141.46             0              0 46 i

Total number of prefixes 10
bgpd> 
```

Validation Result
Final(origin, path)

Indicates the status
of this update

Local Preference
variable (+-) or fixed

Update Identifier

QuaggaSRx (3)

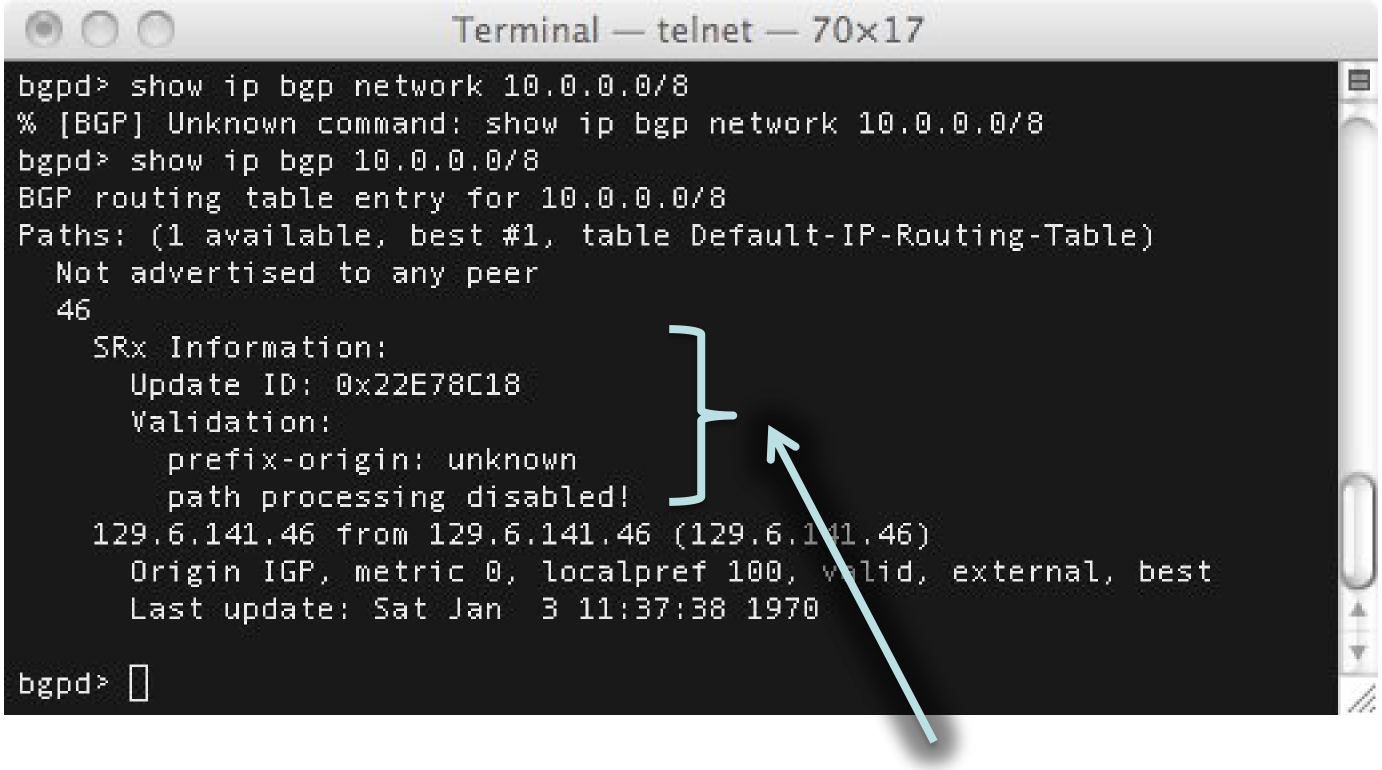

```
bgpd> show ip bgp network 10.0.0.0/8
% [BGP] Unknown command: show ip bgp network 10.0.0.0/8
bgpd> show ip bgp 10.0.0.0/8
BGP routing table entry for 10.0.0.0/8
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  46
    SRx Information:
      Update ID: 0x22E78C18
      Validation:
        prefix-origin: unknown
        path processing disabled!
    129.6.141.46 from 129.6.141.46 (129.6.141.46)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Last update: Sat Jan  3 11:37:38 1970

bgpd> 
```

BGP-SRx Information embedded in
BGP network information