Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking

http://www.caida.org/funding/hijacks/

Summary

Recent reports have highlighted incidents of massive **Internet traffic interception executed by re-routing BGP paths across the globe** (affecting banks, governments, entire network service providers, etc.). The potential impact of these attacks can range from massive eavesdropping to identity spoofing or selective content modification.

Because of their complex dynamics, and the number of different actors involved on a global scale, devising effective methodologies for the **detection and characterization of traffic interception events requires empirical and timely data** (e.g., acquired while the event is still ongoing). Such data must be a combination of **passive BGP measurements and active measurements (such as traceroutes)**, since the mechanism triggering the attack operates on the inter-domain routing control plane, but the actual impact is only verifiable in the data plane.

In this project we:

1. investigate, develop, and experimentally evaluate **novel methodologies to automatically detect traffic interception** events and to characterize their extent, frequency, and impact;
2. extend our measurement infrastructure to **detect in near-realtime and report episodes of traffic interception based on BGP hijacking**;
3. **document such events**, providing datasets to researchers as well as informing operators, emergency-response teams, law-enforcement agencies, and policy makers.

Approach Overview

The detection process is made of **two phases**:

In **Phase A**, we select prefixes and ASes involved in one of the following suspicious **control plane** events:

i. ***Multiple Origin AS(MOAS)*** - a prefix starts being originated by multiple ASes, or ***Increase in prefixes originated by each AS*** - an AS starts originating a larger number of prefixes

ii. ***Violations of the "valley-free" assumption*** - a customer AS offers to transit traffic directed to a specific prefix between its providers

iii. ***New edges in the AS graph*** - a BGP announcement contains an adjacency that has never been noticed in the topology before

iv. ***Inconsistent path prepending*** - the number of prepended ASes is changed by a third AS in the path

PHASE A			
MOAS Prefix count increase per AS	> no interception > impersonation	**1 > multiple origin interception**	> black hole hijack
Valley free violation	**2 > interception via valley free violation** > special agreements > misconfiguration	> suspicious data plane and control plane mismatch	> misconfiguration > black hole hijack
New edge	> new connection > impersonation	**3 > interception via path poisoning**	> misconfiguration > black hole hijack
Inconsistent prepend	**4 > interception via prepend manipulation**	> suspicious data plane and control plane mismatch	> misconfiguration
PHASE B	AS paths match	AS paths are different	Could not reach prefix

In **Phase B**, for each potential victim prefix we analyze the results of the **data plane** measurement and we compare the resulting AS path with the one observed in the control plane.

If the AS paths match, and either ***event (ii)*** or ***event (iv)*** were observed in the control plane, then we are observing an interception via valley-free violation (2) or an interception via prepend manipulation (4).

If the AS paths differ, and either ***event (i)*** or ***event (iii)*** were observed in the control plane, then we are observing a multiple-origin interception (1) or an interception via path poisoning (3).

Attack Scenarios

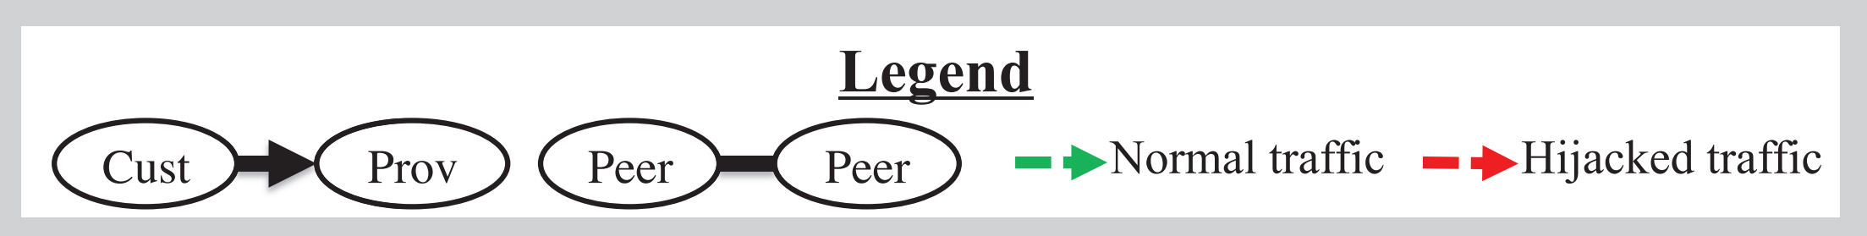

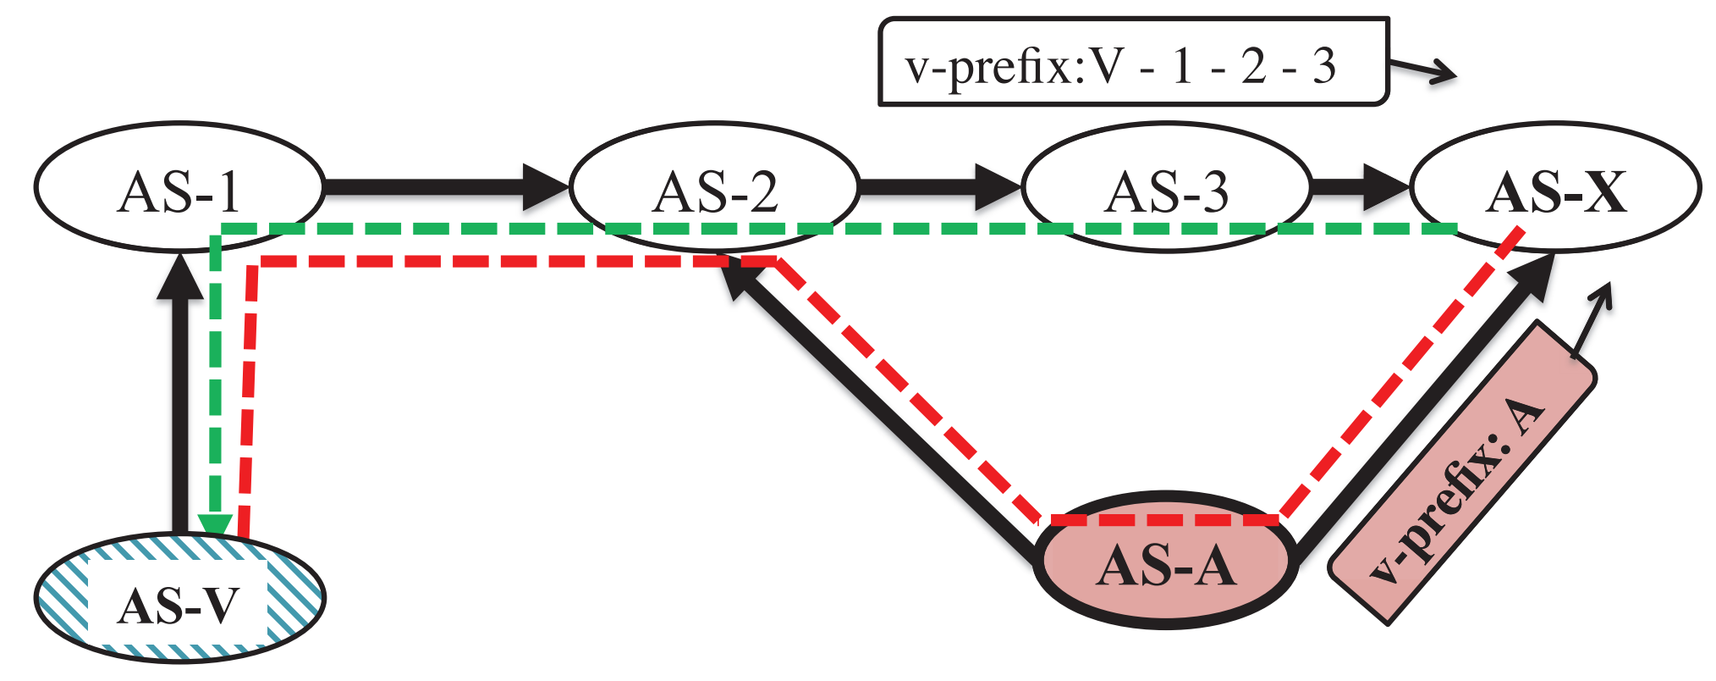

1. Multiple origin interception - the attacker AS announces the victim prefix and then redirects the traffic through the original route

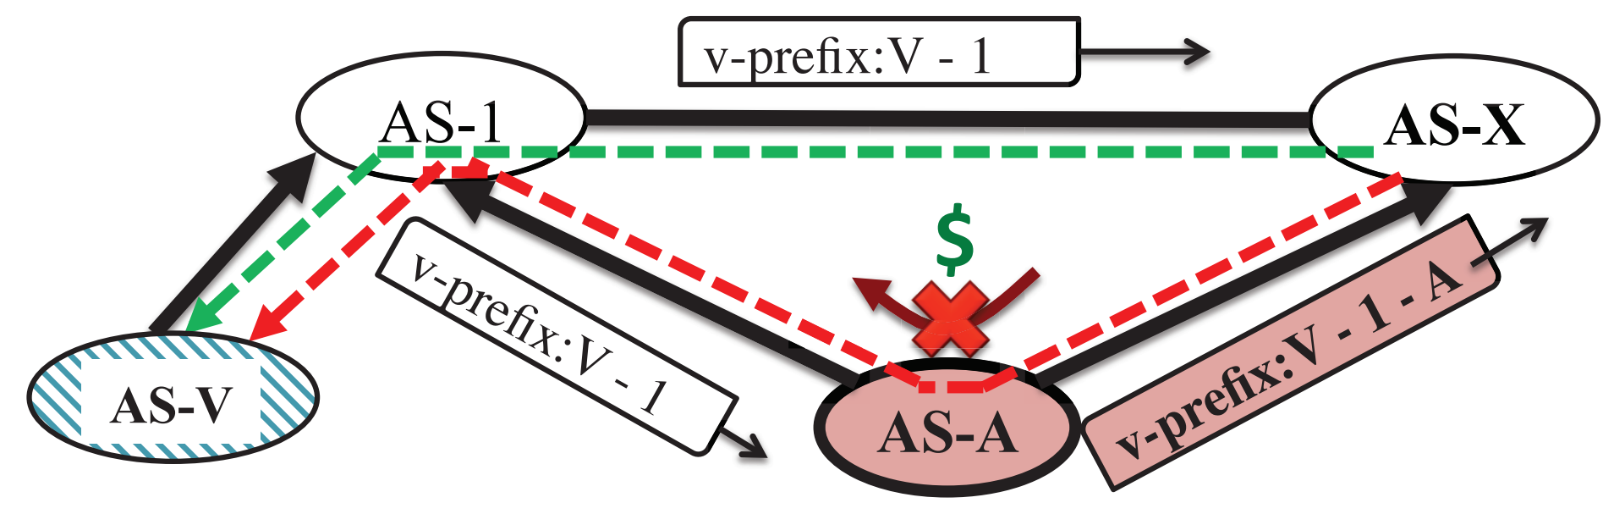

2. "Valley-free"violation - the attacker AS transits traffic between its providers, thus intercepting data that would usually not route

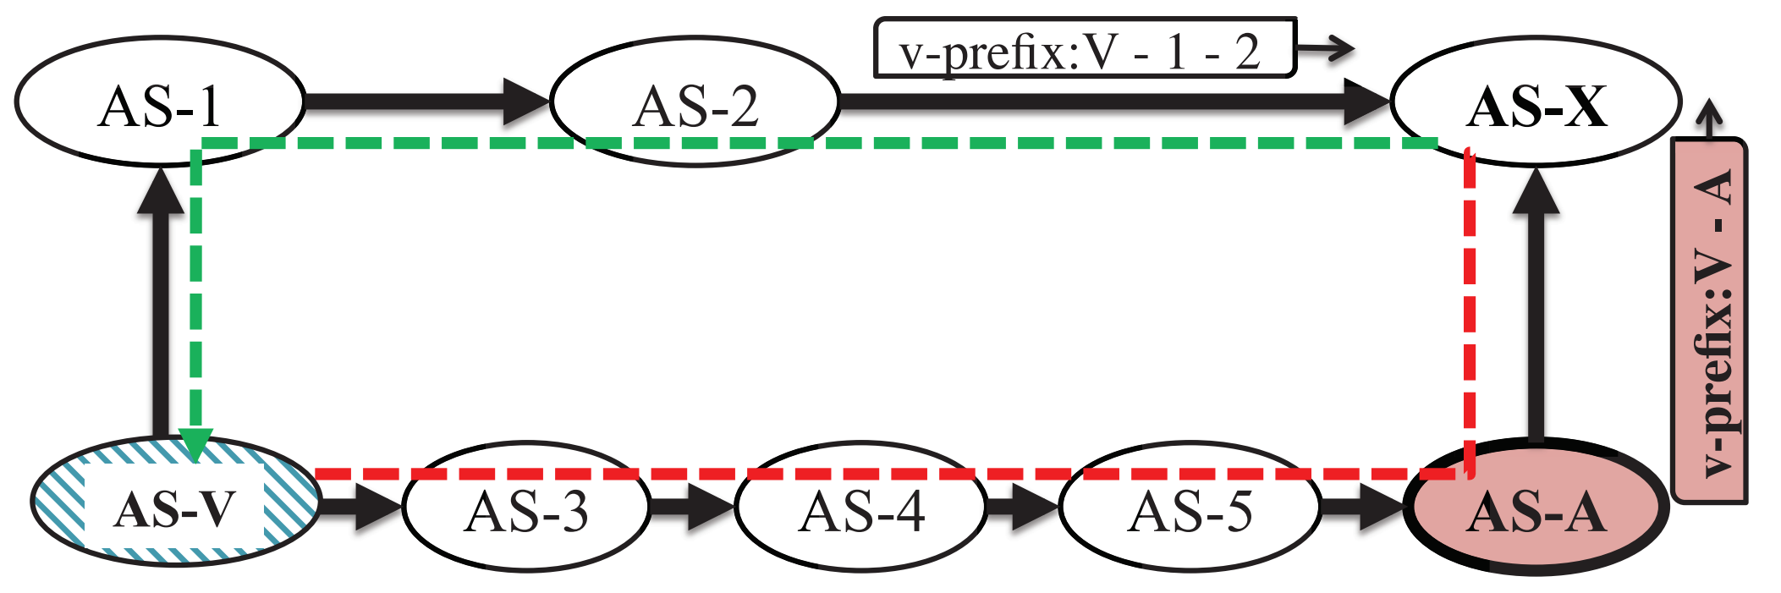

3. Path poisoning interception - the attacker AS announces false adjacencies in its path toward the victim AS and attracts its traffic

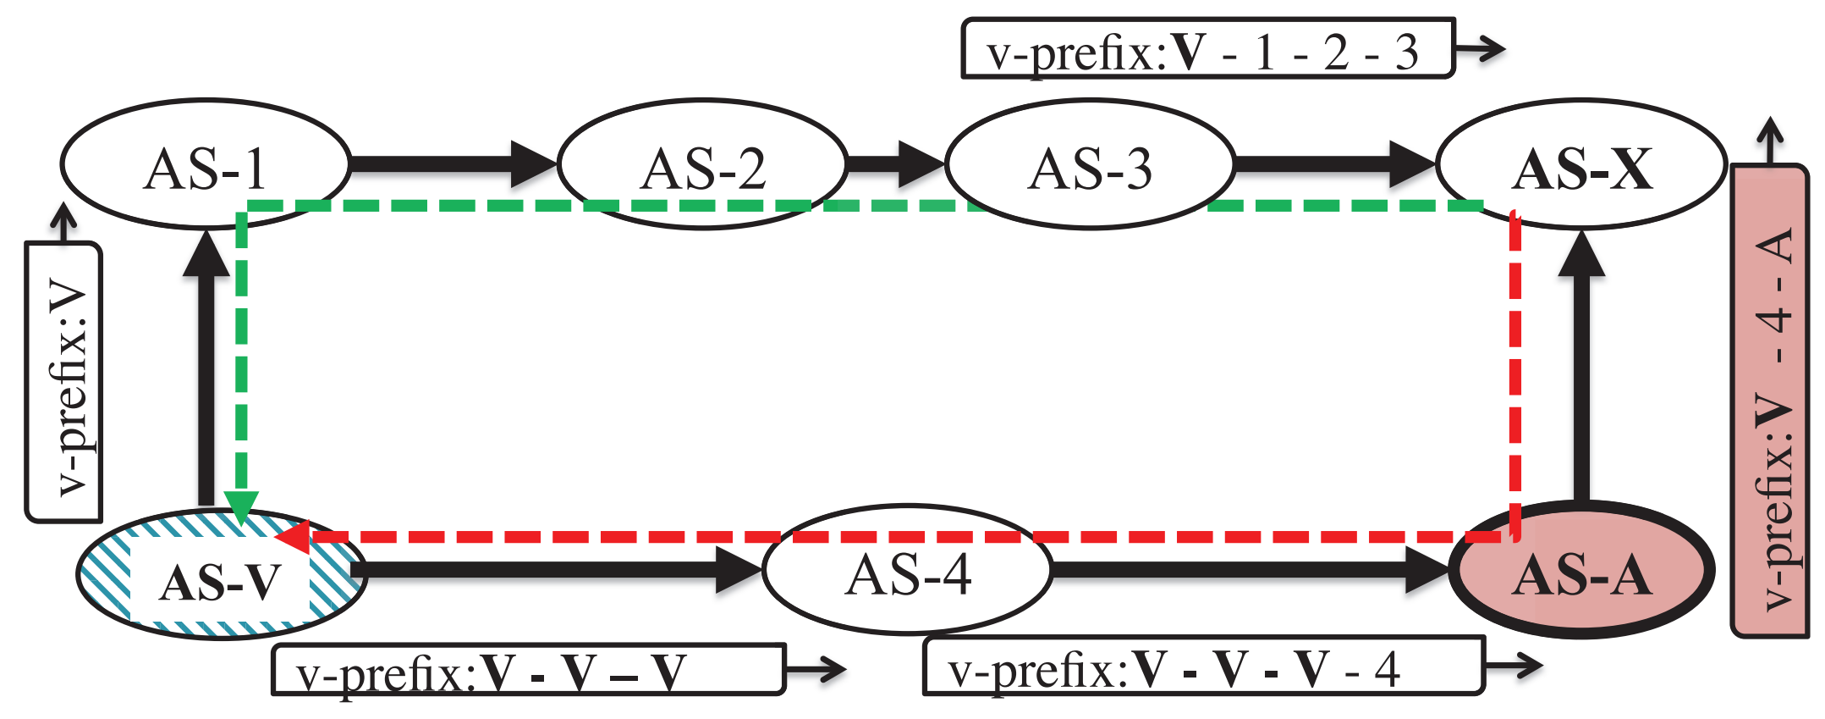

4. Prepending manipulation - the attacker AS removes the prepended list of ASes that have been added by other AS to alter the route that traffic takes

Architecture Overview and Timeline

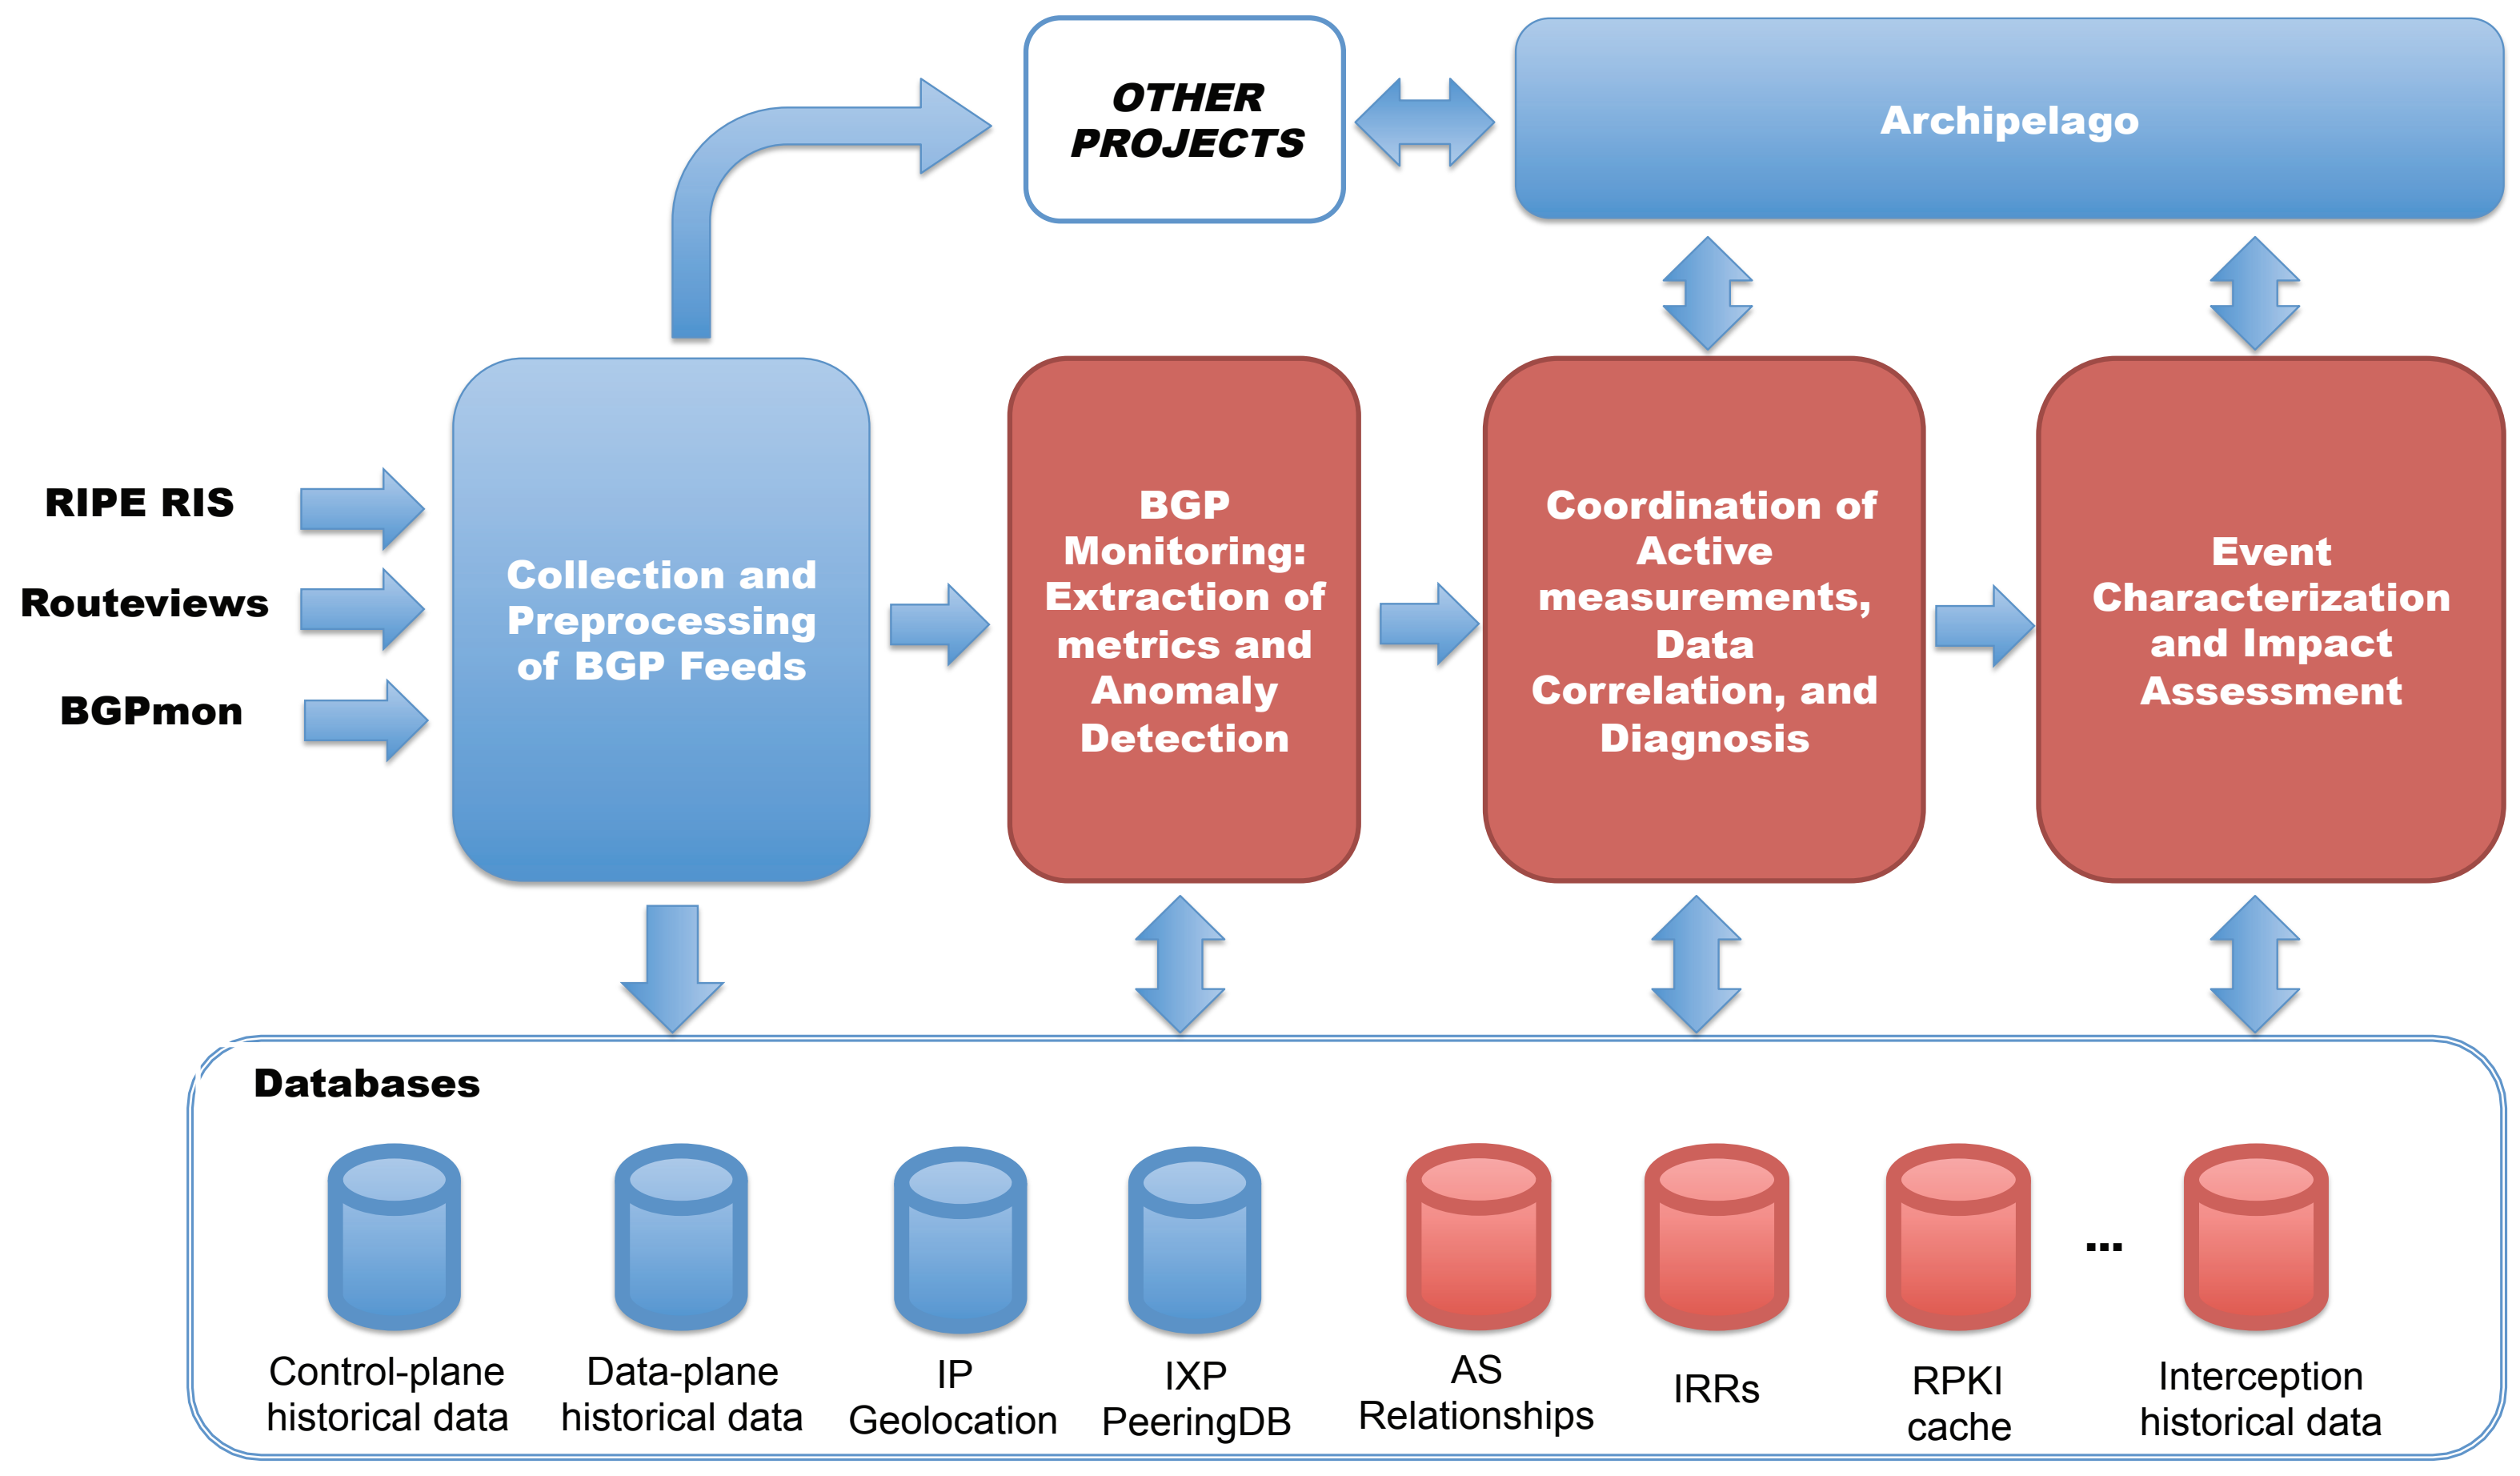

OTHER PROJECTS
Archipelago
RIPE RIS
Routeviews
BGPmon
Collection and Preprocessing of BGP Feeds
BGP Monitoring: Extraction of metrics and Anomaly Detection
Coordination of Active measurements, Data Correlation, and Diagnosis
Event Characterization and Impact Assessment
Databases
Control-plane historical data
Data-plane historical data
IP Geolocation
IXP PeeringDB
AS Relationships
IRRs
RPKI cache
...
Interception historical data

Our infrastructure for data collection and analysis. Portions in red denote components to be developed and integrated as part of this proposal, whereas components deployed within previous or ongoing synergistic projects are in blue.

Geographical distribution of Archipelago nodes

Task 1 : Extend and refine infrastructure for data collection and analysis (Years 1, 2, and 3);

Task 2 : Design a method for detecting and characterizing traffic interception (will start in the second half of Year 1);

Task 3 : Collect and disseminate data and knowledge: organize a workshop, blog on incidents, engage operators, provide data to researchears (Years 2 and 3).

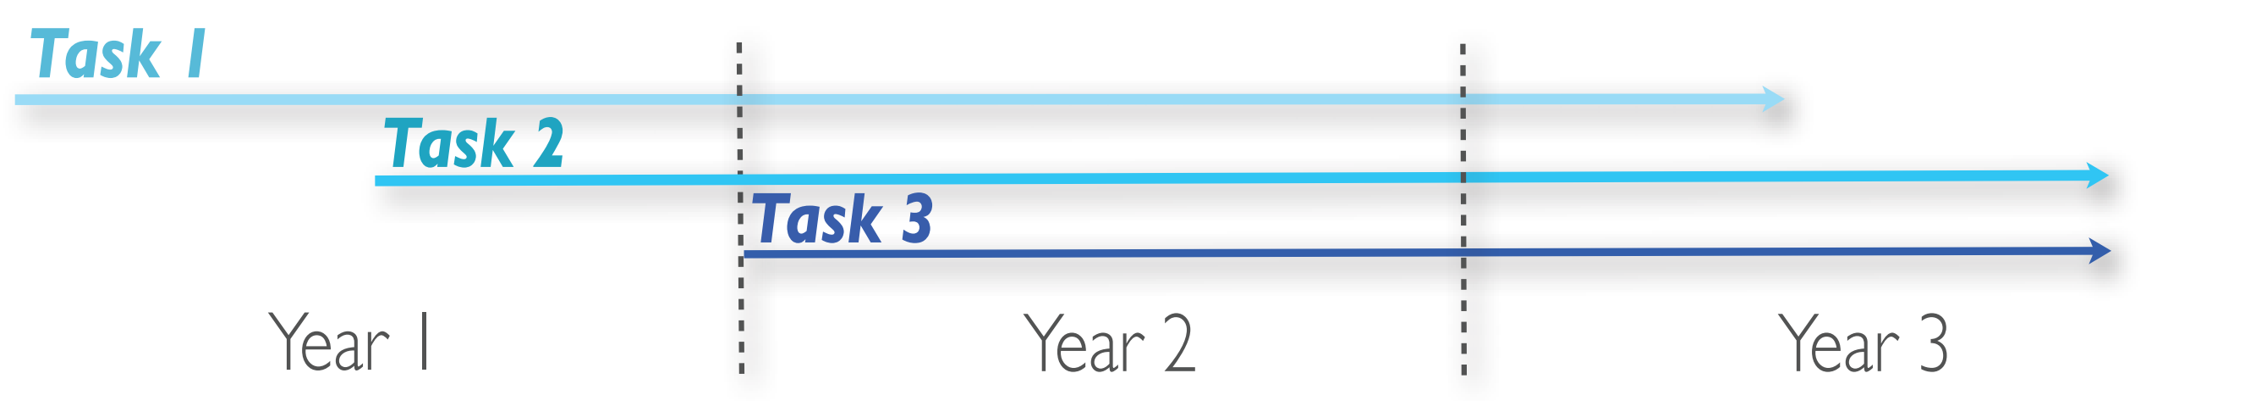

January 2015

team
Alberto Dainotti | Phillipa Gill | Chiara Orsini
Alistair King | Vasco Asturiano | kc claffy

sponsored by
Funding source:
NSF CNS-1423659.

UC San Diego
Stony Brook University
SDSC SAN DIEGO SUPERCOMPUTER CENTER
caida