Carrier Grade IP?

Albert Greenberg
Jennifer Yates
Fred True
AT&T Labs-Research

Agenda

Why carrier grade IP?

What makes it hard?

Solution approach and roadmap

Focus

- Information Systems for fault/performance data management
  - Payoffs in improved network and service

Why Carrier Grade IP?

Increasing number of diverse applications over IP

- Data, Web, Voice, Video (IPTV), Gaming, …

Increasingly stringent requirements

- Commerce / business critical transactions
  - Outages expose enterprises to huge losses
- Web-based apps “24x7”
  - Activity at all hours – when to schedule maintenance?
- Performance sensitive applications
  - Small network glitches cascade that trigger large application outages

Increasing pressure to scale

- More service, more infrastructure, lower cost, fewer people

What Are We Up Against?

Hard, long lasting failures:

- Fiber cuts, router failures, line card failures, …
- Hardware and *software* problems
- Approach: design and control for diversity/resilience; engineer net mgt systems for rapid service restoration

Chronic, intermittent faults:

- Outages that “clear” themselves, but keep recurring
  - Impact that adds up, even if the per event impact is small
- Hardware and *software* problems
- Approach: engineer net mgt systems for forensics + network/systems update to prevent recurrence

Solution and Roadmap

Removal of single points of failure

Fast and reliable failure detection

Fast service recovery (restoration)

Fast fault repair

Hitless maintenance

What about the edge?

- Cost and single points of failure concentrated at the edge
  - Innovation: 1:N interface sparing, 1:N router sparing (router farm)

What about the network (edge + core)?

- Fast diagnosis for real time response and off-line forensics
  - Innovation: network data management systems that simplify analysis of complex and/or massive network data

Focus: Information systems for fault/performance data management -- Goals

Scale: Efficient storage of potentially large and complex data feeds over long periods of time

Feature-Richness: Comprehensive capabilities for data querying and reporting, which could be used to construct a variety of higher level applications

Speed: Support for real-time data

Ease of Operation: Very low maintenance and management overhead: "DBA-less"! (DBA = Data Base Administrator)

- Straightforward paradigm for adding new feeds/tables: "Wizard-like"
- Automatic creation of various database mechanisms: bulk data ingest, load control scripts, schema, data aging, logging/alerts
- Automatic configuration of logs, alerts

Open design: Employ the use of "open" toolsets where possible

What's Hard in Network Data Management?

Data Distribution – Getting data where it needs to be without complex, disjoint interfaces

- **Solution:** Data Distribution Bus

Managing Change – Constant churn of new data, changing record layouts and schema, field values, etc.

- **Solution:** Automation and code generation

Keeping Track of Things – Managing a coherent catalog of metadata: loading status, schema, business intelligence (e.g. field validation rules)

- **Solution:** Integrated metadata database and query tools

Scale – Building a system with features that scale evenly. Harnessing parallelism throughout the design; scaling "outside the box".

- **Solution:** Daytona™ data management system – provides scale, stability, speed – optimized for reliable processing of reliable data

Maintaining Uniformity across Data Sources – facilitating data correlation and combining; encouraging the use of common conventions, field types, keys, etc.

- **Solution:** Automation brings homogeneity to the data model!

Logical Data Architecture

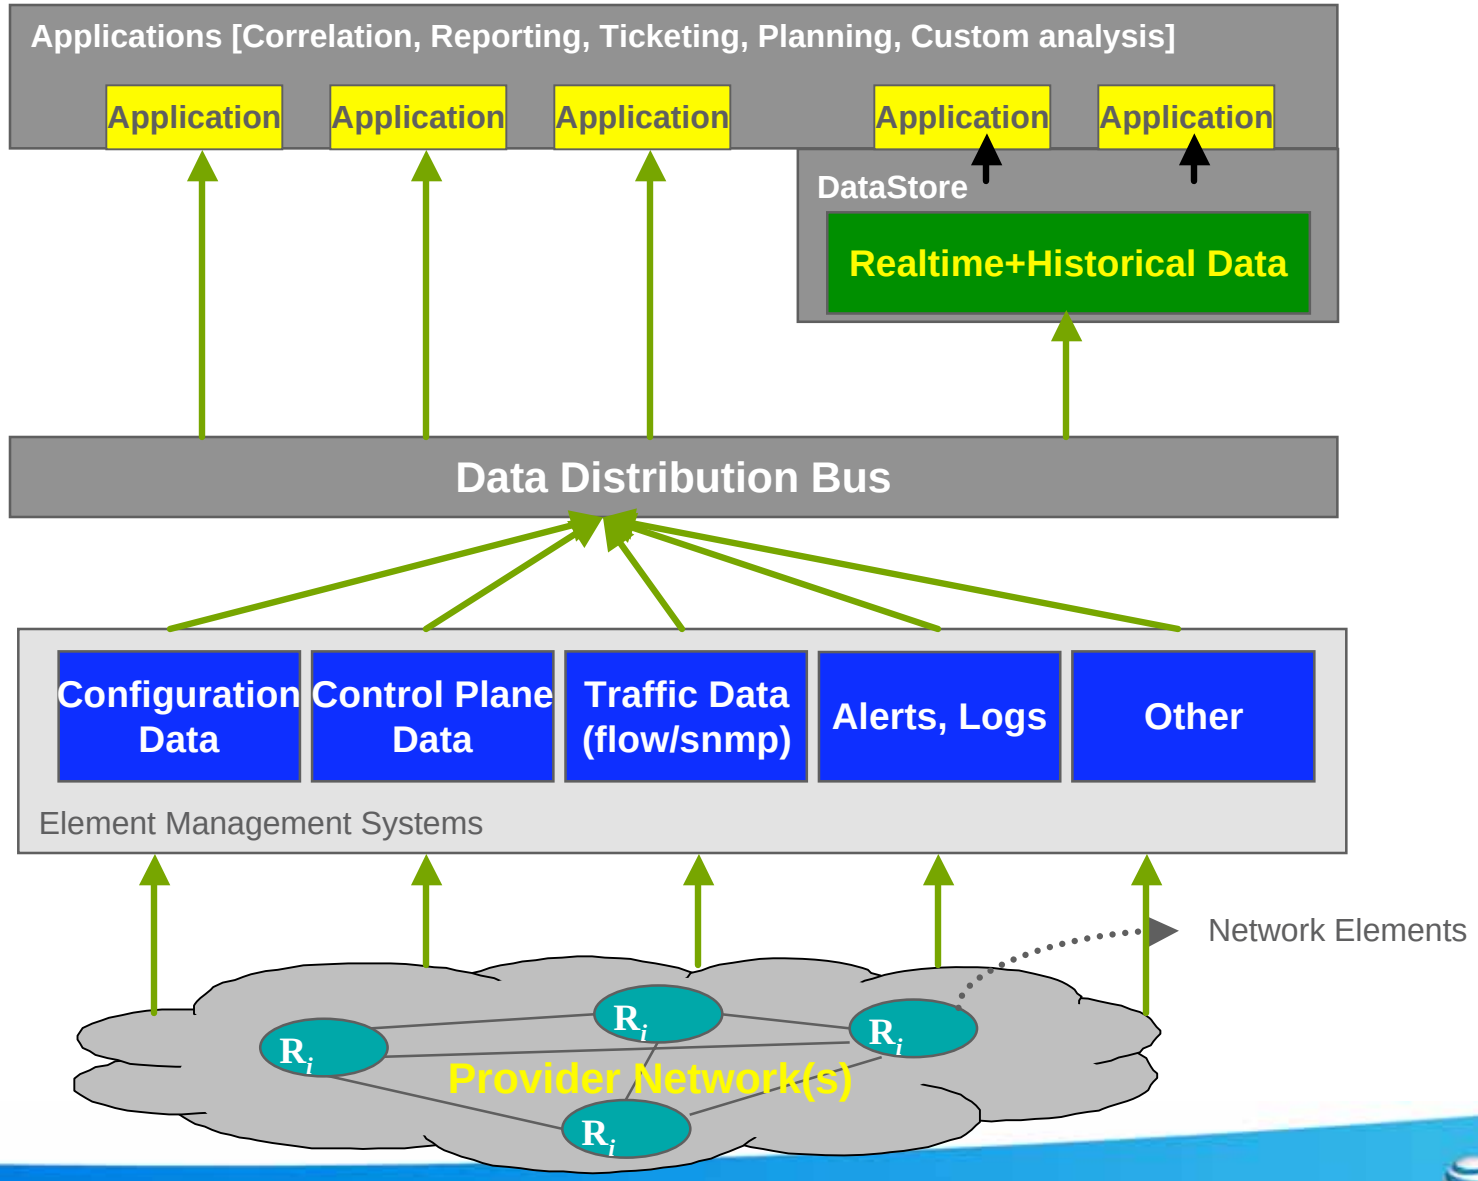

Data Distribution Bus

- Data Distribution Bus: “Glue” between data collectors, repositories, and reporting/analysis systems. **One** logical system (and associated business process) to transport data everywhere.

- Automated data transit management, data tracking, and “publisher/subscriber” model
- Decoupling of data publishers from data subscribers: easier to manage
- 1-time configuration for each publisher/subscriber
- Inherently parallel/scalable
- Unified interface for all publishers/subscribers
- Unified alerting/alarming
- Short-term recovery buffer for critical data

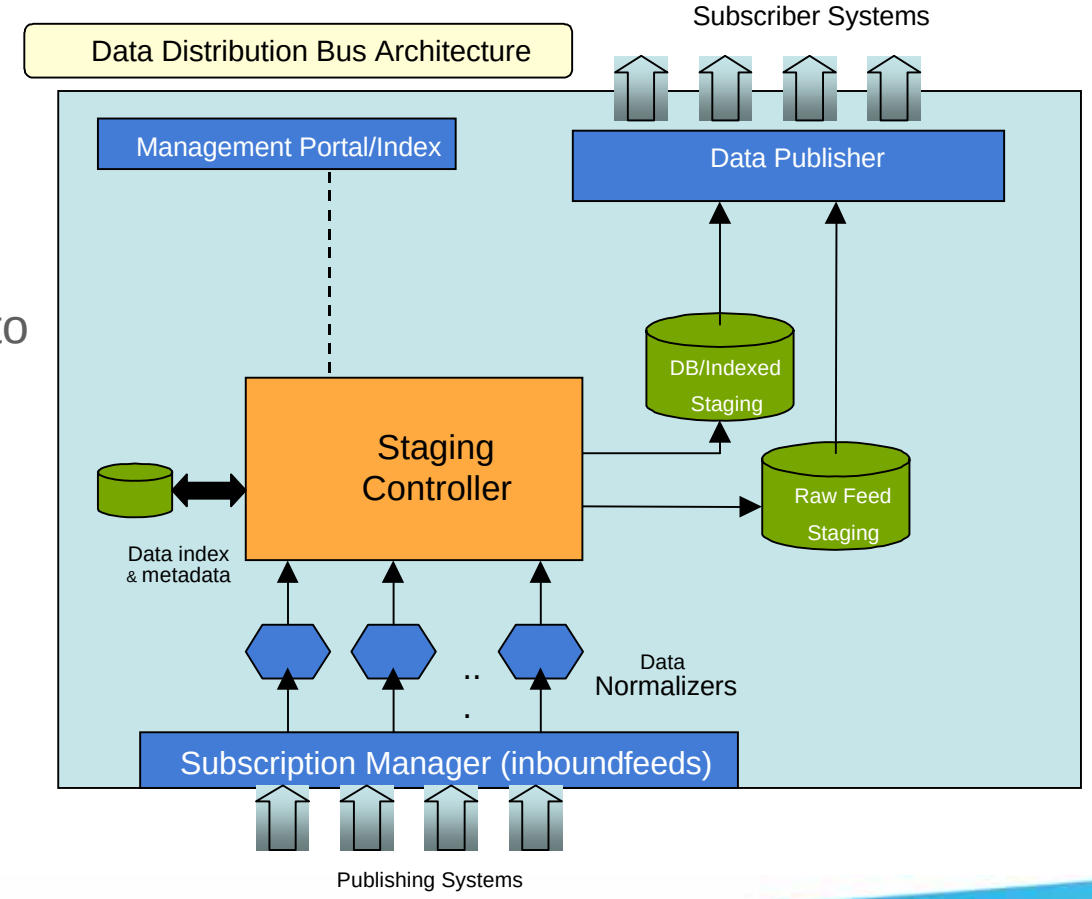

Data Automation

Systems for services

- Ingest
- Validation
- Logging
- Versioning
- Retry
- Aging of old data
- …

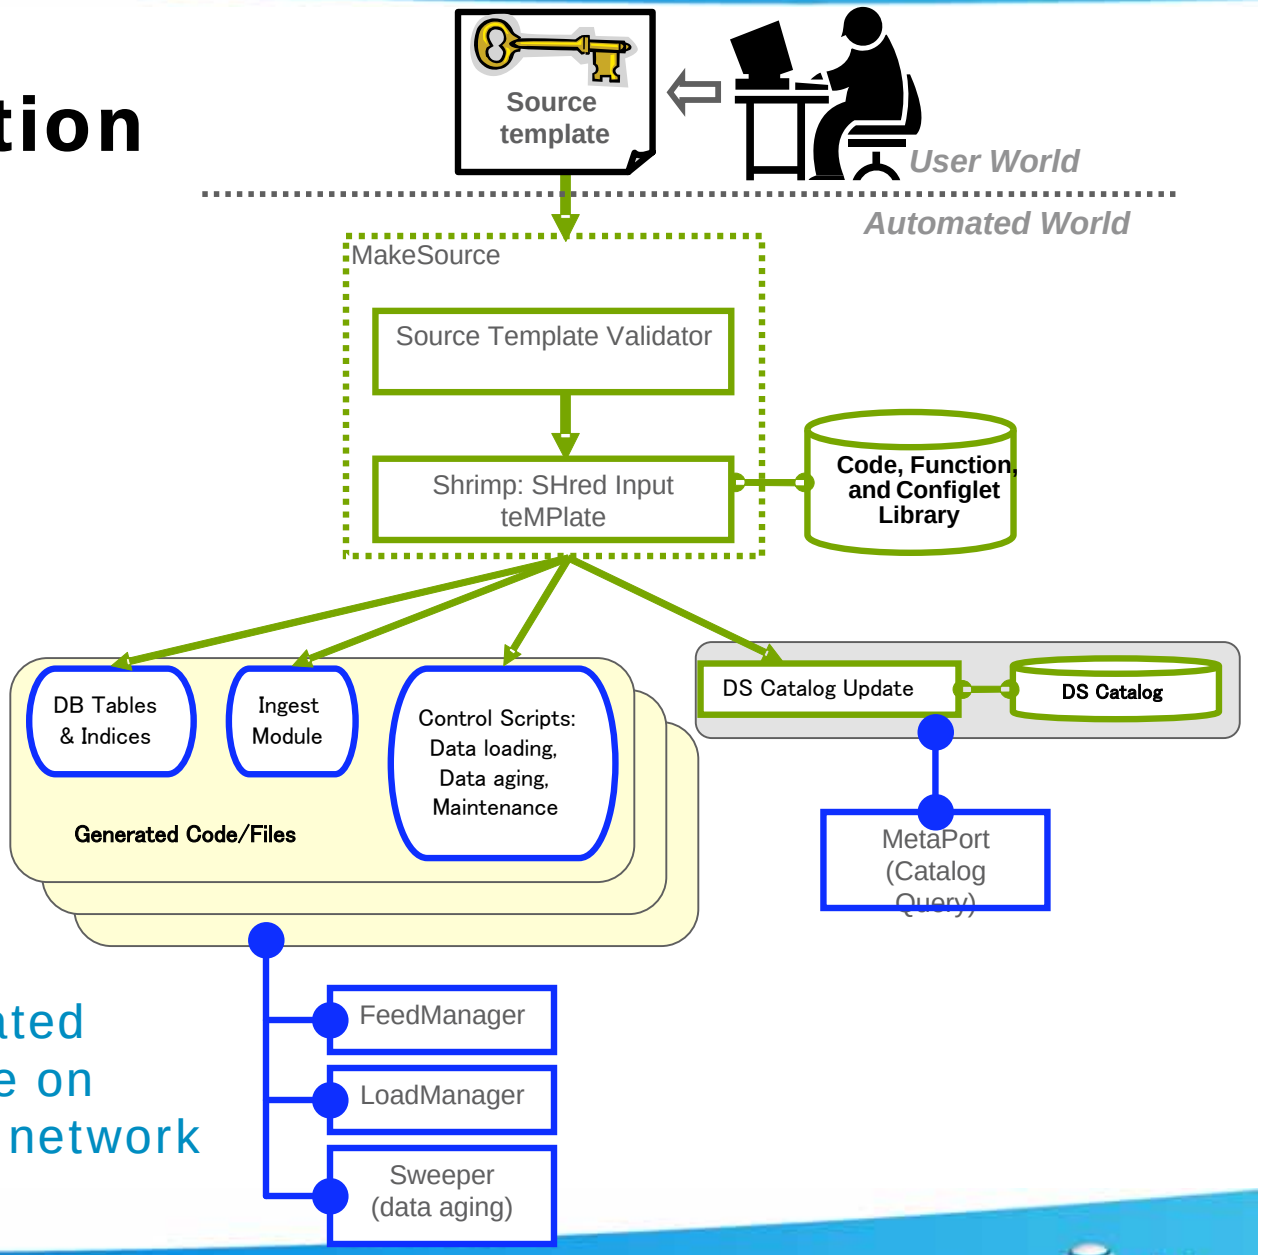

So that OPs and automated systems can concentrate on *challenges* in improving network reliability

Automated Analytic Toolkit

Libraries for temporal, spatial clustering

Pairwise network data time series correlation testing

Chronic, intermittent fault identification (temporal correlation)

Silent fault localization (spatial correlation)

Reduced false alarm rate

Automated, rapid classification of all performance impacting events

Real time and offline customer trouble shooting

- Select edge interfaces, network paths, traffic, applications by customer
- Select customer traffic, services, applications by network element

Fault prediction

Example: CPU Anomalies & Link Load

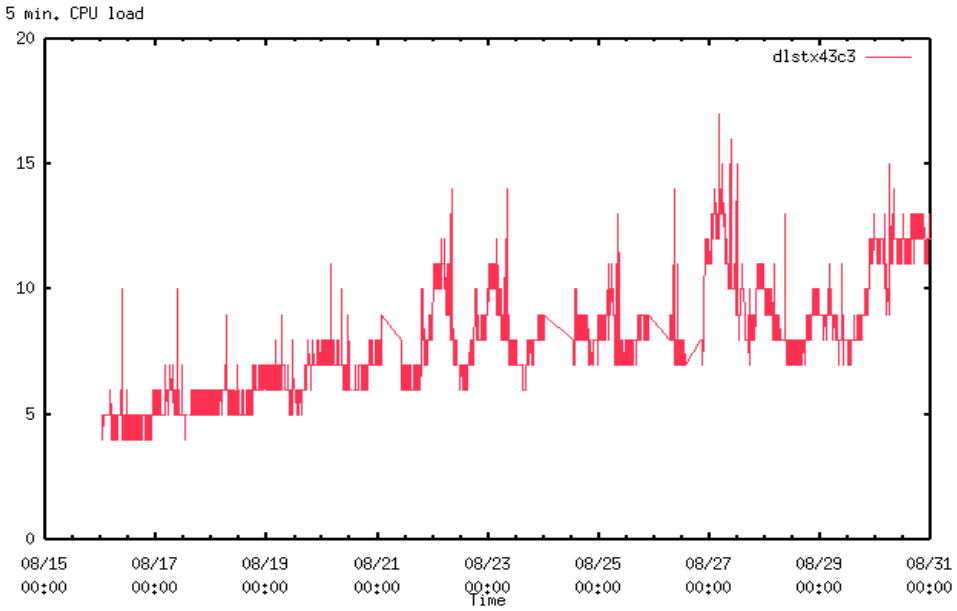

Anomaly detection identifies unusual behaviour; Correlation testing identified routers with CPU varying daily with load – surprising!!!!

- Some observed increasing over time

Operations' forensics tracks this down to subtle configuration issue

- Closed out a DOS vulnerability, potentially amplifying small attacks on an interface into total router failure
- Automated global configuration repair

at&t

Example: Cross-Layer and Automated Correlation

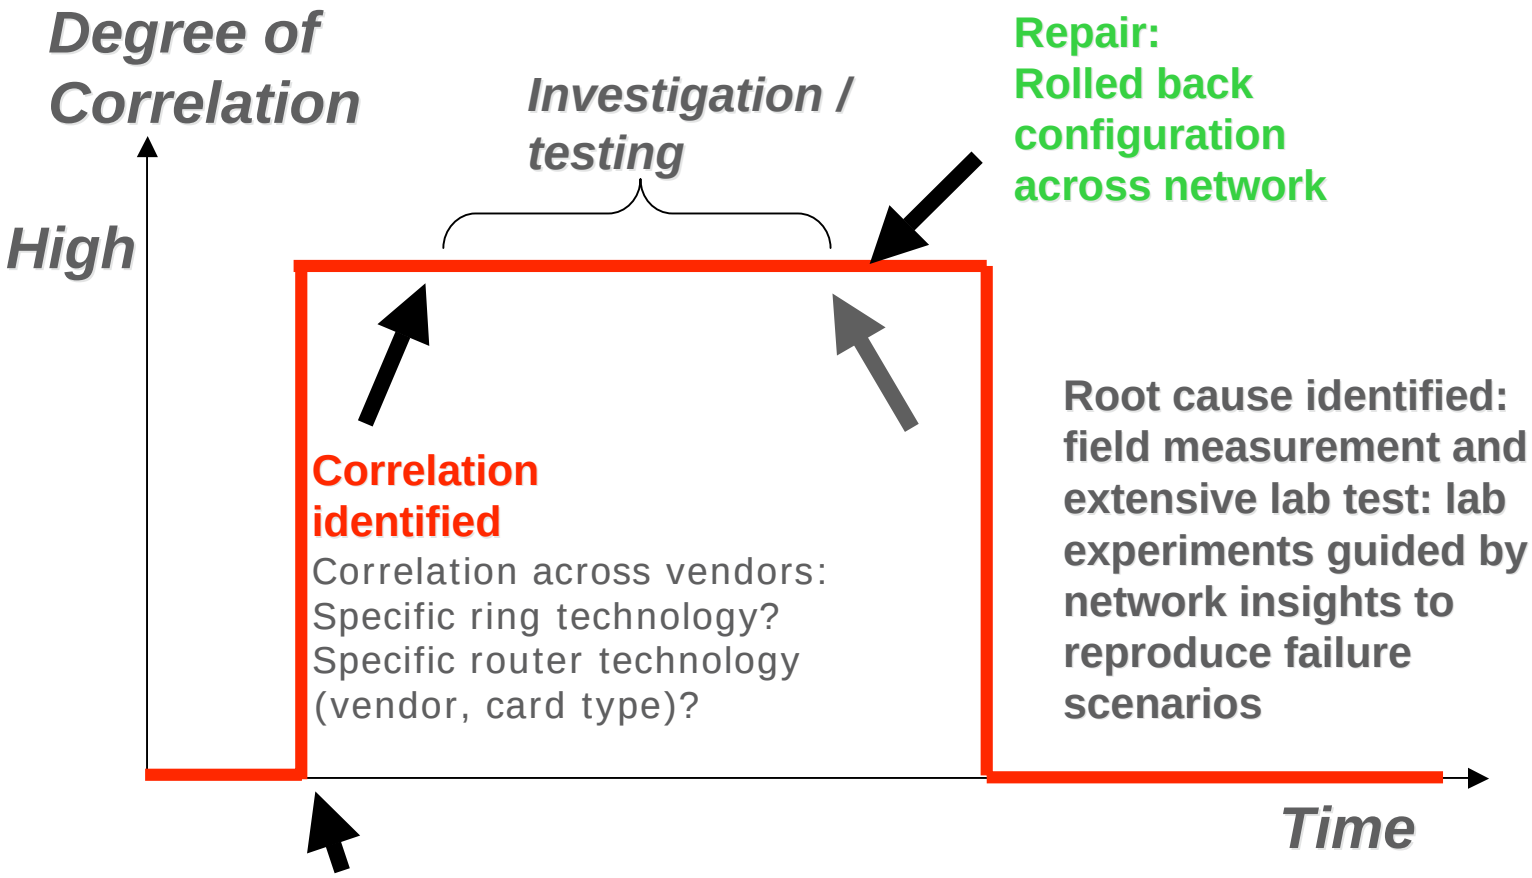

SONET Rings and IP Link Failures

at&t

Example: Silent Failure Localization

Real time localization of outages for rapid failure recovery

- Particularly for “silent” faults (i.e., no alarms generated to indicate which network element is having a problem)
- Designed to operate in harsh network environment
  - Multiple simultaneous failures
  - Missing data
- Correlate end to end monitoring alerts with topology to find most likely fault location

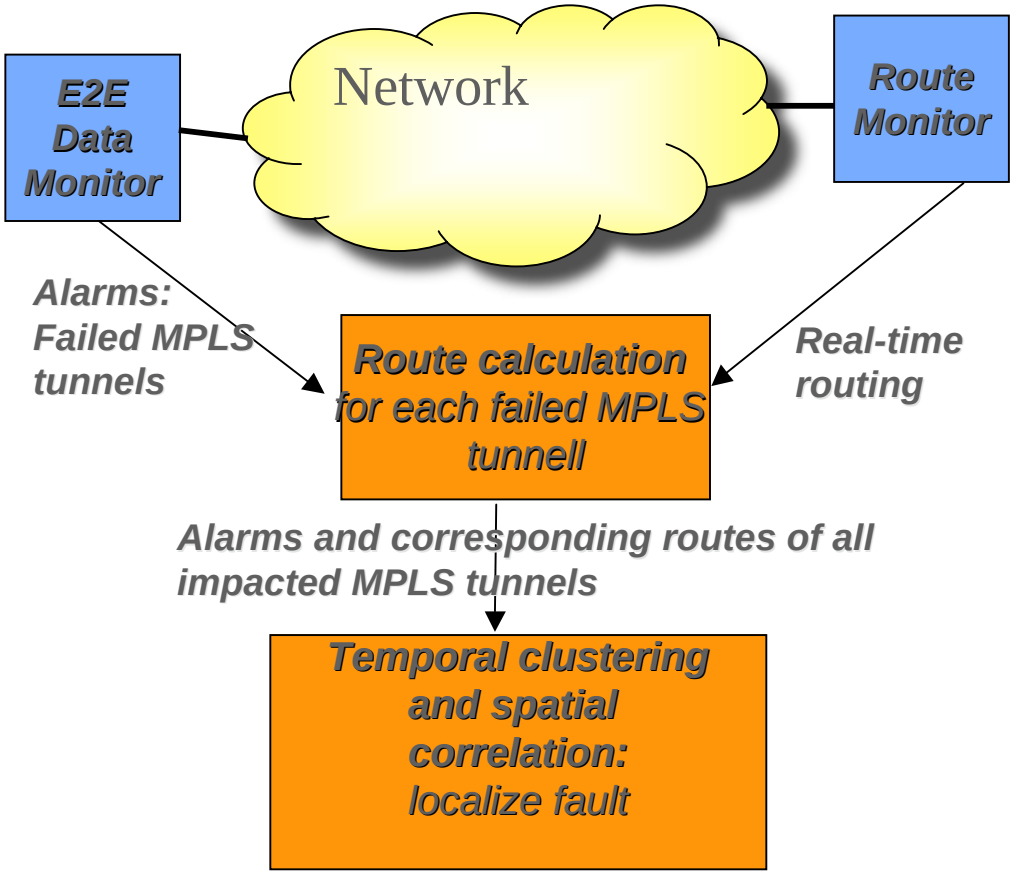

at&t

Outcomes

Improved Network

- Identification and permanent removal of egregious problems, that had been flying under the radar

Improved Network Management Systems and Processes

- Faster service restoration and network repair

Detect → Localize → Diagnose → Restore / Repair

Corollary benefits: codify policy/configuration, self-documenting network, share knowledge across organizations

How To Push Automation As Far As Possible

Timely, accurate information is essential!!!

- Example: Precise topology and available capacity now

Tools that separate capabilities from policies, since policies can change fast

- Example: link utilizations should be < 80% except for links involved in that new VoIP trial in Phoenix with vendor X equipment, where utilizations should be < 50%, except for …

Statistical versus those rooted in domain expertise?

Big Guard Rails – extensive monitoring and info correlation/validation

Huge Operations involvement at every step

- Simpler, repeatable tasks/repairs automated