IRATI
INVESTIGATING RINA

Pristine

The Pouzin Society

Welcome to the RINAissance!

An Introduction
to the RINA Architecture
Part 2

IRATI RINA Workshop

John Day

Dublin 2014

In a network of devices why would
we route between processes?
- Toni Stoey, RRG 2009

What a Layer Looks Like

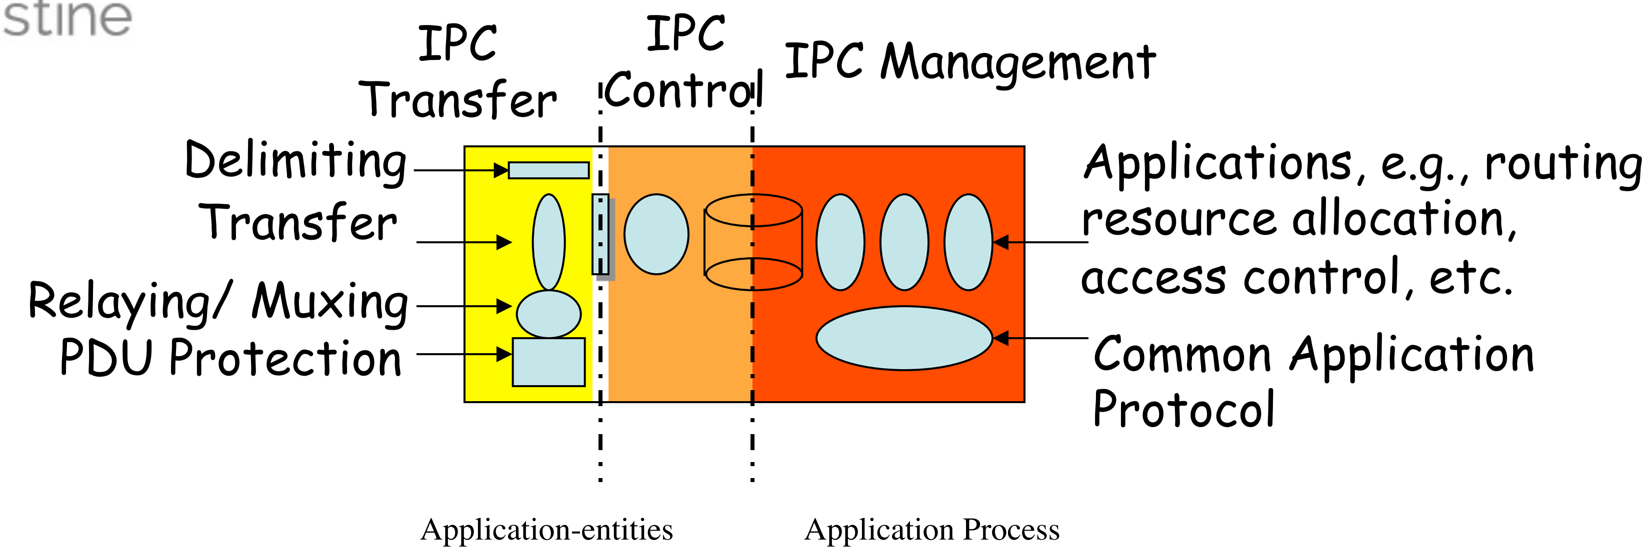

- Processing at 3 timescales, decoupled by either a State Vector or a Resource Information Base
  - IPC Transfer actually moves the data (≈ IP + UDP)
  - IPC Control (optional) for retransmission (ack) and flow control, etc.
  - IPC Layer Management for routing, resource allocation, locating applications, access control, monitoring lower layer, etc.
- Remember that within a scope if there is a partitioning of functions, it will be orthogonal? Well, here it is.

What are the Protocols?

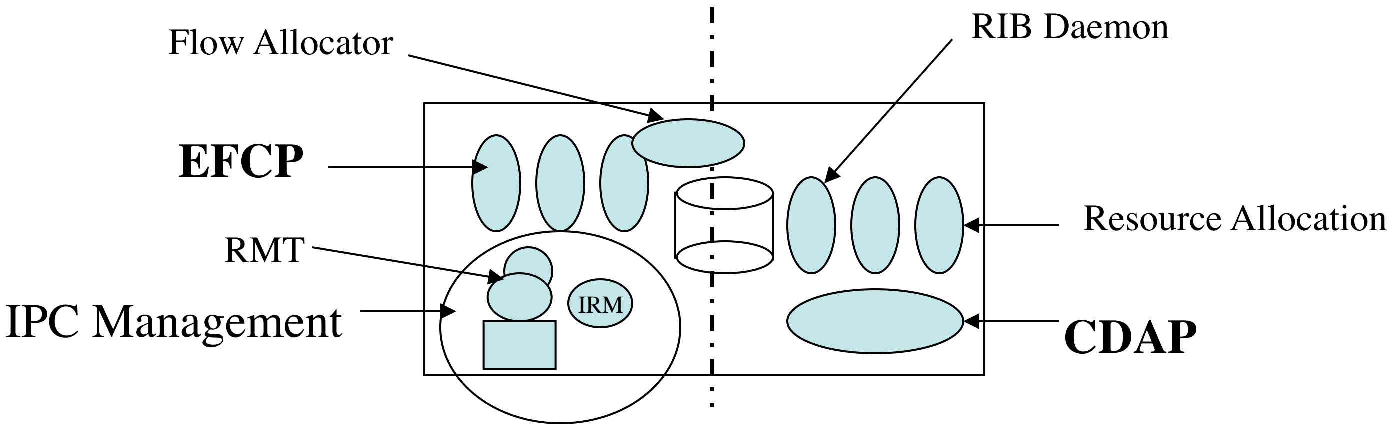

- Only two
  - A data transfer protocol, **EFCP**, based on delta-t with mechanism and policy separated. This provides both unreliable and reliable flows.
  - The common application protocol based on **CDAP**.

Error and Flow Control Protocol* (EFCP)

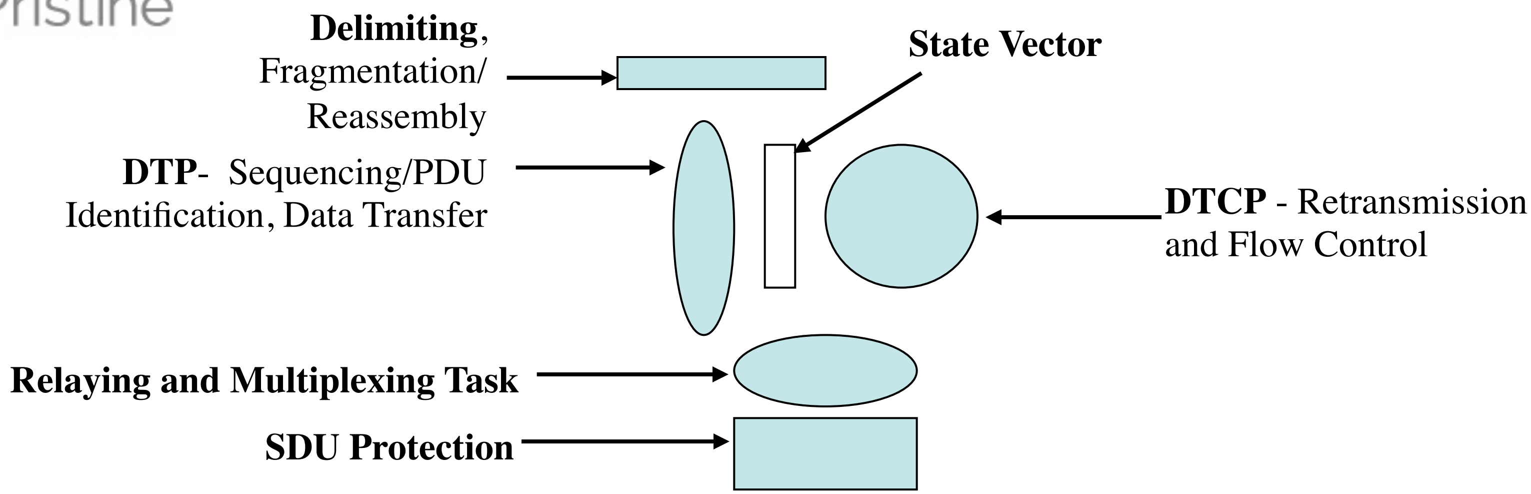

- There is one DTP (and possibly one DTCP) per flow
- One Relaying and Multiplexing Task per IPC Process
- At most, one SDU Protection per (N-1)-flow/DIF.
- Delimiting is a function with an inverse that converts SDUs to User-Data Fields for PDUs.
- DTP is a bit like IP/UDP (if you don't look too close)

- See the EFCP Specification

*A Bit More to Say About

Error and Flow Control Protocols

- This has been a subject of intense research since the early 70s. You would think we had pretty much exhausted the subject. Not so. First there was:
  - Separating mechanism and policy revealed that the protocol naturally cleaved into data and control and
  - that there was no distinct relaying protocol and
  - SDU Protection wasn't really part of the protocol and
  - Heavy weight solutions like IPsec were unnecessary
- Recently we have also found that:
  - Fragmentation/Reassembly is part of Delimiting
  - The A-Timer has a role in DTP.
  - Finer grained flow control across connections is possible external to EFCP.
  - Can dynamically shift from using DTCP without affecting a connection?
    - Very likely. Set DRF and just do it.
- Why are we finding these new results?
  - Because we are acting scientifically: we are trying to understand the nature of the problem.

IRATI INVESTIGATING RINA

Common Distributed Application Protocol

The Pouzin Society

Pristine

- Common Application Connection Establishment (CACE) is the common procedure for establishing application connections. It ensures that there is a known first exchange.
  - Based on the OSI ACSE which was defined to be used recursively and has hooks for. . .
- An authentication module, which is a policy of whatever strength required.
- CACE and an authentication module can be wrapped around any existing application protocol, e.g. HTTP. (See the CACE specification.)
- CDAP provides the minimal six operations and a basic object-oriented functionality scope and filter. (See the CDAP specification.)
  - Would other programming paradigms lead to different functions?

Only Three Kinds of Systems

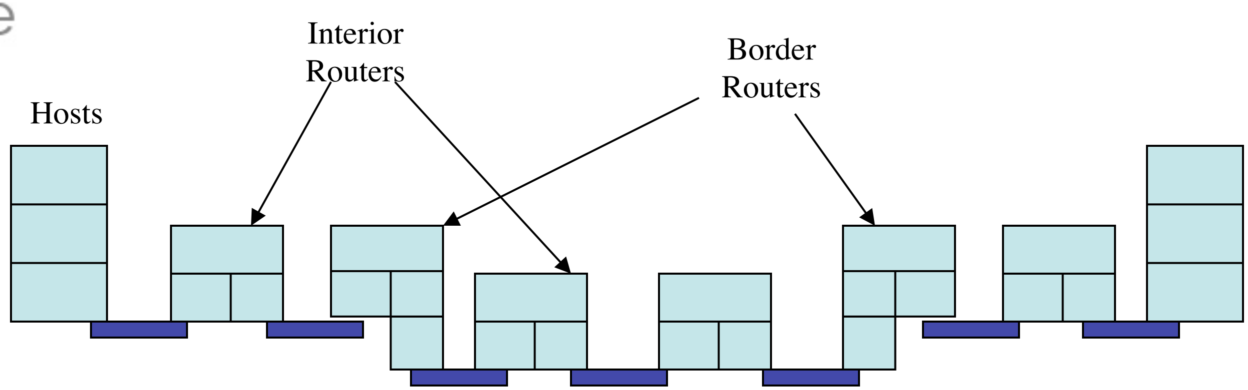

- Middleboxes? We don't need no stinking middleboxes!
- NATs: either no where or everywhere,
  - NATs only break broken architectures
- The *Architecture* may have more layers, but no *box* need have more than the usual complement.
  - Hosts may have more layers, depending on what they do.

Hosts Might Have More DIFs

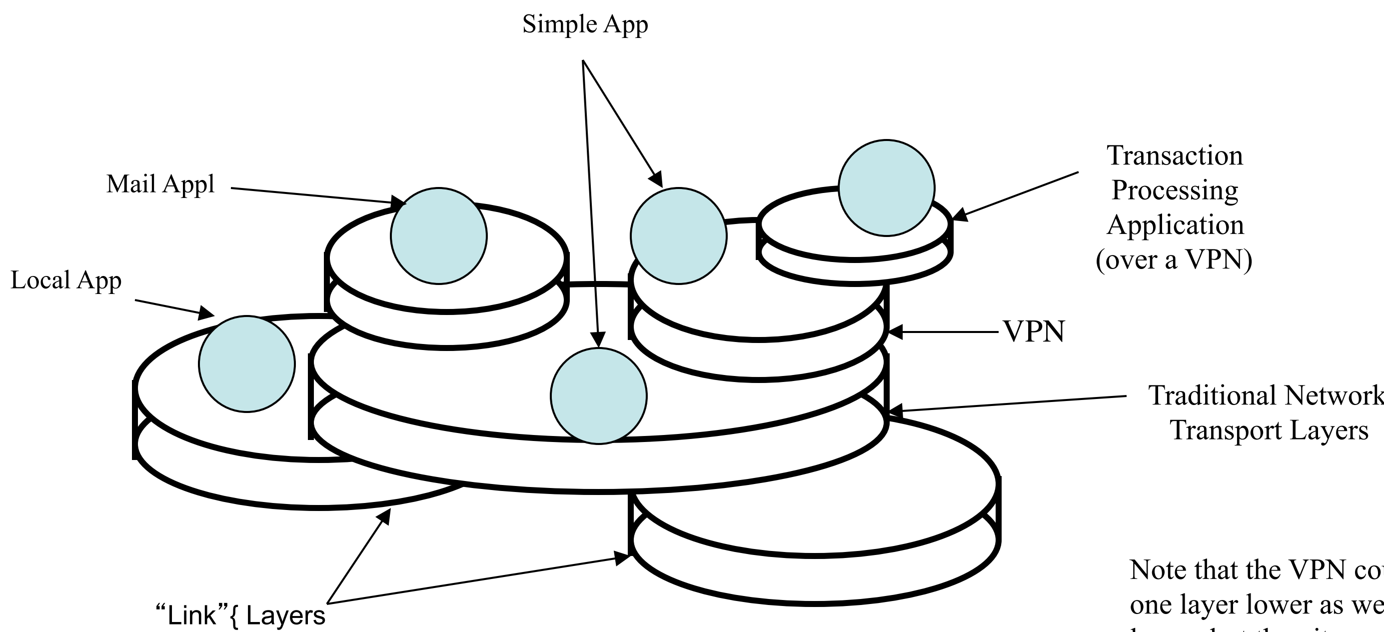

Note that the VPN could occur one layer lower as well or even lower, but then it would just be a PN.

User Applications use whatever layer has sufficient scope to communicate with their apposite.

All Communication goes through Three Phases

- Enrollment
  - Operations to create sufficient state within the network to allow an instance of communication to be created.
- Allocation (also known as Establishment)
  - Operations required to allocate an instance of communication creating sufficient shared state among instances to support the functions of the data transfer phase.
- Data Transfer
  - Operations to provide the actual transfer of data and functions which support it.
- Most of our attention has been on the last two. The first has often been ignored and is usually seen as necessarily ad-hoc. But enrollment turns out to be key.

How Does It Work?

Enrollment or Joining a Layer

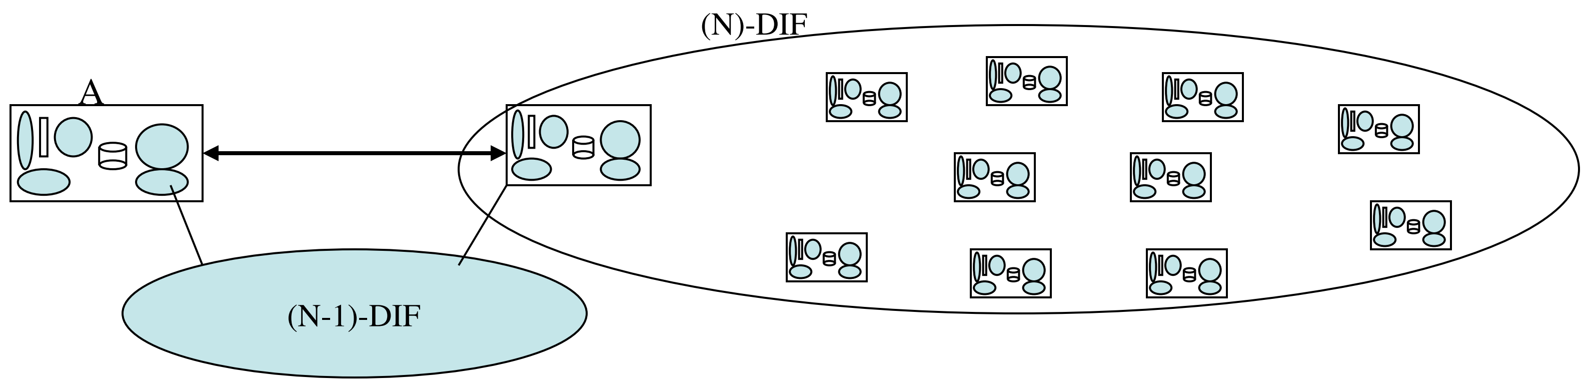

- Nothing more than Applications establishing communication (for management)
  - Authenticating that A is a valid member of the (N)-DIF
  - Initializing it with the current information on the DIF
  - Assigning it a synonym to facilitate finding IPC Processes in the DIF, i.e. an address
  - (see the Enrollment specification for an example.)

Enrollment

Details: The Naïve Approach

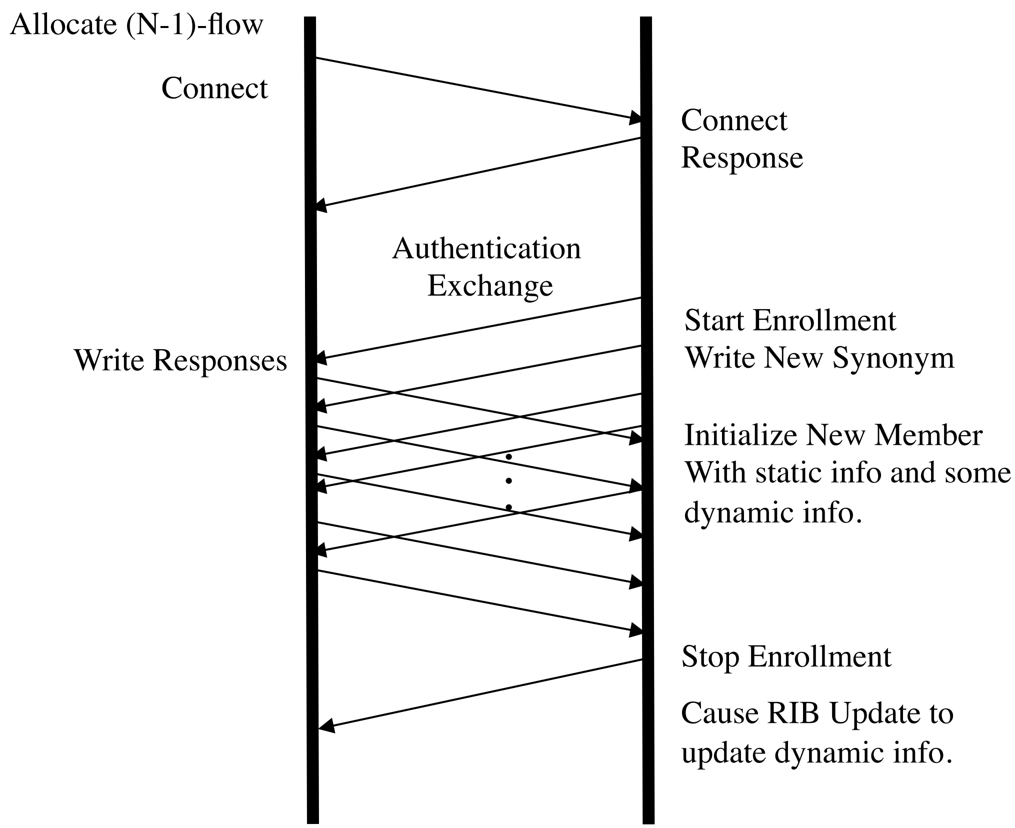

- First, create a flow with the lower layer.
- Then, create an application connection with a member of the DIF at this layer, then authenticate each other. If successful, proceed to enroll.

IRATI
INVESTIGATING RINA

Pristine

Enrollment

Resource Allocation Layers

The Pouzin Society

- If a member fails and reappears (router crash), can't assume all is well.
  - OTOH, the naïve approach would be onerous.
- Assume that most Enrollment requests are members dropping out and coming back.
- Assume that state information they have (including addresses assigned to them) has a lifetime. If gone a short amount of time, then assume it is still good; otherwise, upload. (But still has to authenticate)
- But to save time, let the new member decide what to upload.

Enrollment

A Pragmatic Approach for Resource Sharing Layers

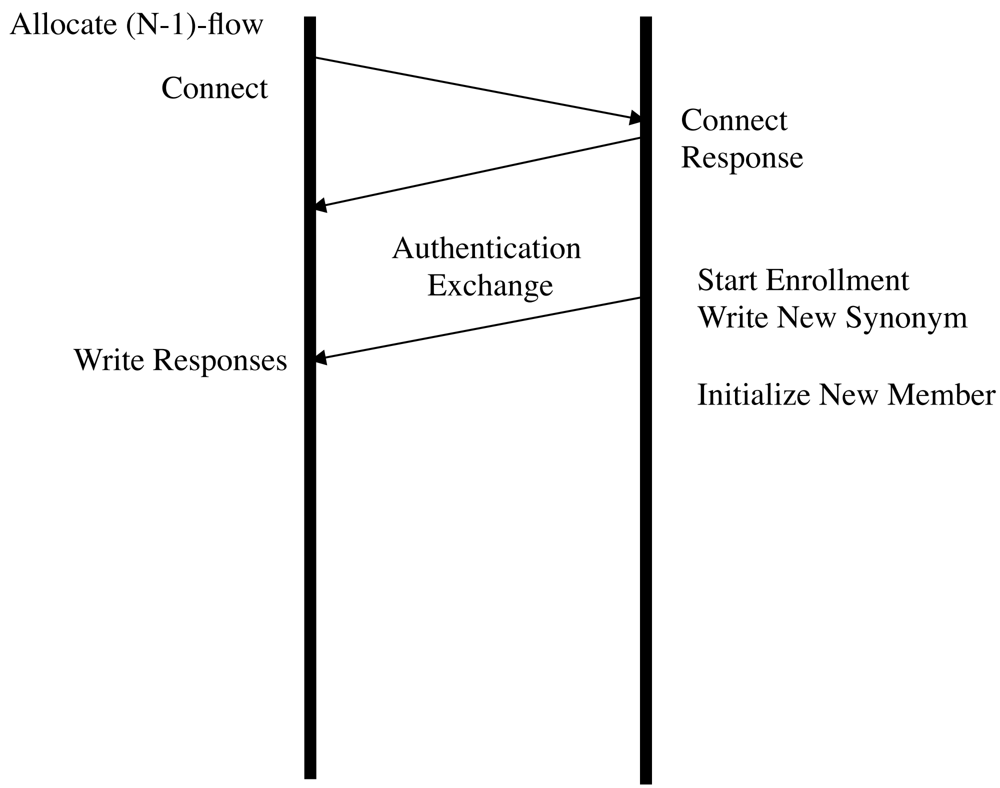

- First part is the same.

Enrollment Procedure I

- *When the New Member receives the M_Connect Response, the New Member copies Current_Address to Saved_Address, it sends*
  - → *M_Start Enrollment(address, Address_expiration_time, other data about New Member)*
- /* The New Member is telling the Existing Member what it knows. Primarily this is derived from the address (NULL or not), and the expiration life-time of the address if non-NULL. Since addresses are generally assigned for hours or minutes, tight time synchronization is not required. (Even for DIFs with fast turnover, fairly long assignment times are still prudent.)*/
- The Member sends
  - ← M_Start_R Enrollment(address (potentially different), Application Process Name, Current_Address, Address_Expiration).

Enrollment Procedure II

- Using the information, provided by the New Member, the Existing Member sends
  - ← M_Create (zero or more) to initialize the Static and Near Static information required. When finished and the New Member has sent all necessary
  - → M_Create_Rs
- The Existing Member sends a
  - ← M_Stop Enrollment (Immediate:Boolean)
- The New Member may Read any additional information not provided by the Existing Member.
  - → M_Read (zero or more)
  - → M_Stop_R Enrollment
- If the Immediate Boolean is True, the New Member is free to transition to the Operational state.
- If the Boolean Immediate is False, then the New Member can not transition to the Operational state until an M_Start Operation is received.

Enrollment Procedure III

- The New Member is free to Read any information not provided by the Existing Member. Once these are completed, the Existing Member sends:
  - ← M_Start Operation
- The New Member sends
  - → M_Start_R Operation
- Invoke RIB Update of dynamic information which will cause others to send data to the New Member.

How Does It Work?

Establishing Communication

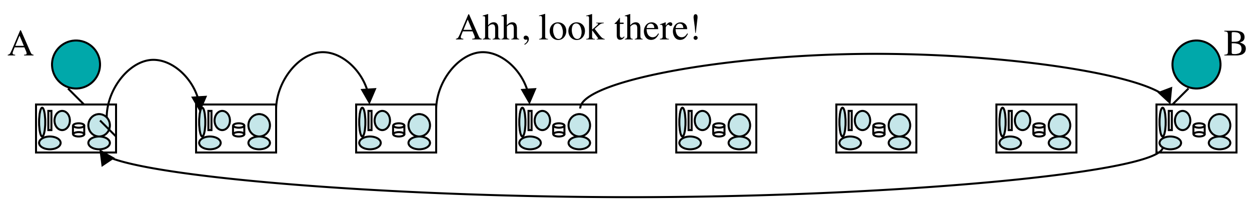

- Simple: do what IPC tells us to do.
  - A asks IPC to allocate comm resources to B
  - Determine that B is not local to A use search rules to find B
  - Keep looking until we find an entry for it.
  - Then go see if it is really there and whether we have access.
  - Then tell A the result.
  - (See Flow Allocator specification)
- This has multiple advantages.
  - We know it is really there.
  - We can enforce access control
  - We can return B's policy and port-id choices
  - If B has moved, we find out and keep searching

Flow Allocator

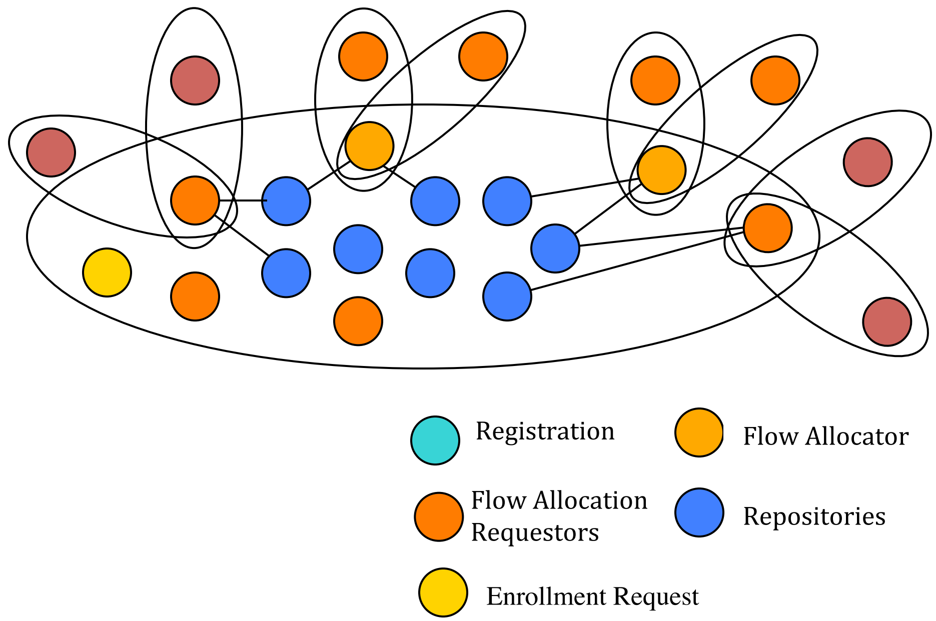

- The Flow Allocator contains a Name Space Management function which registers applications, allocates new addresses within the DIF, and allocates flows.
- Most IPC Processes will maintain a cache of recent name to address mappings from flow allocate requests. Larger DIFs may dedicate IPC Processes to repositories (partially replicated databases) to facilitate flow allocation.

"Internet" Congestion Control

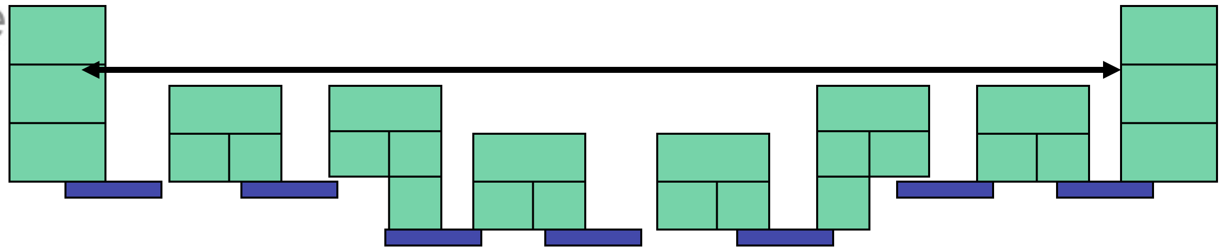

- Congestion Control has been a known issue since 1972.
  - Except in the Internet who only discovered it when it crashed around their ears in 86
- The effectiveness of any congestion control is directly related to the time to effect a change.
  - The longer it takes the less effective the congestion control
- End-to-end implicit notification is predatory.
  - Longest response time. Will work against any attempt to do it at a lower level with shorter scope and better response time.
- The Internet has *network* congestion control,
  - not *internet* congestion control

What is the Difference Between Flow Control and Congestion Control?

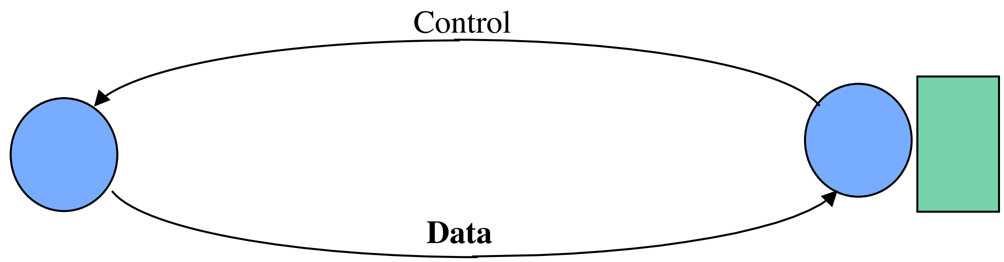

Flow Control is a feedback mechanism co-located with the resource being controlled.

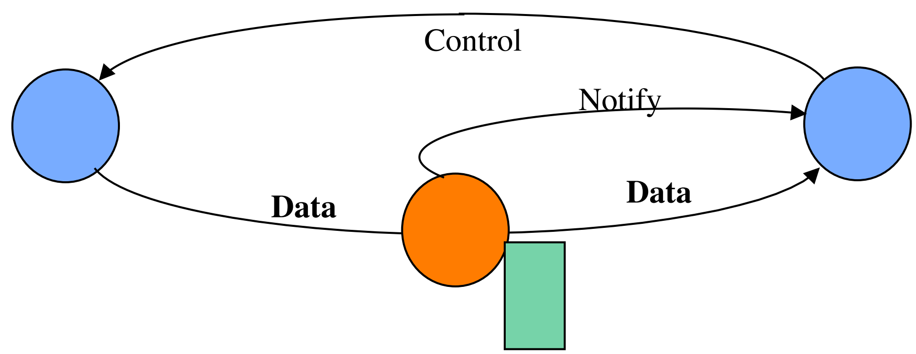

Congestion Control is a feedback mechanism *not* co-located with the resource being controlled.

How Does It Work?

"Congestion Control"

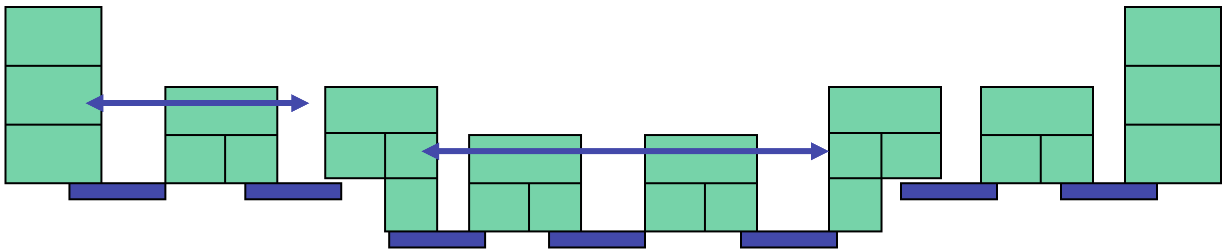

- Congestion Control in TCP was always known to be a stop-gap.
- A DIF always has the potential for the full capability of functions.
- Do flow control (without retransmissions) between intermediate points.
  - Better congestion control, really flow control
  - Allocate different resources to different e-malls.
  - Allows provider much more effective management of resources.
  - Provides means to throttle flows being used for denial of service attacks
  - All of these places? Probably not all in the same DIF. Major Area for Research

How Does It Work?

The Internet and ISPs

- ISPs have as many layers as they need to best manage their resources.

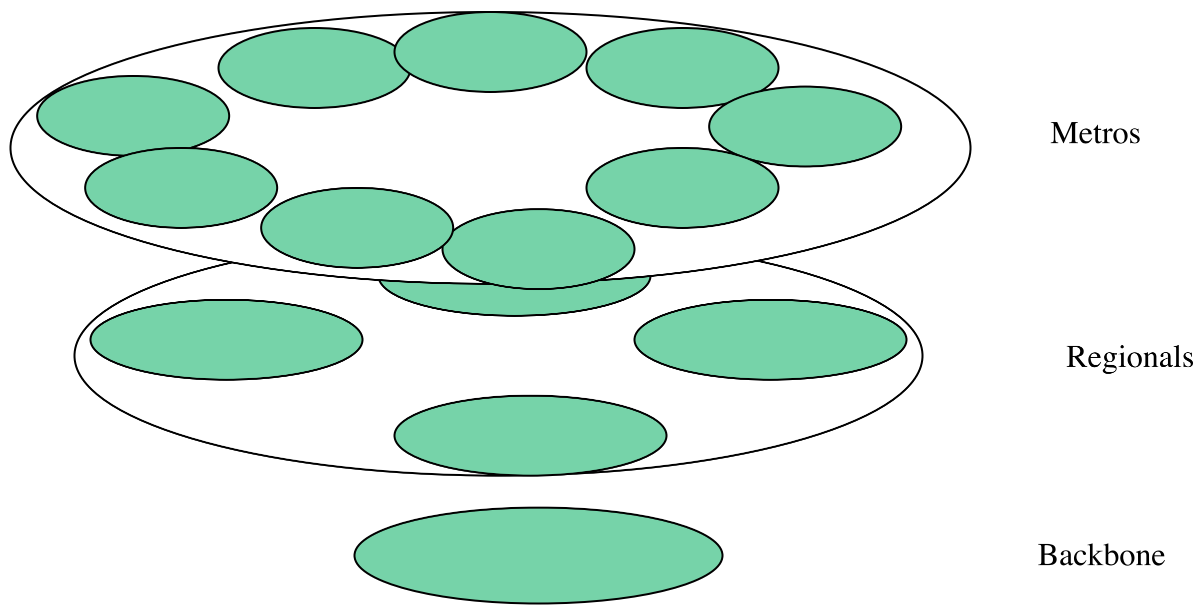

How Does It Work?

The Internet and ISPs

- The Internet floats on top of ISPs, a "e-mall."
  - One in the seedy part of town, but an "e-mall"
  - Not the only emall and not one you always have to be connected to.

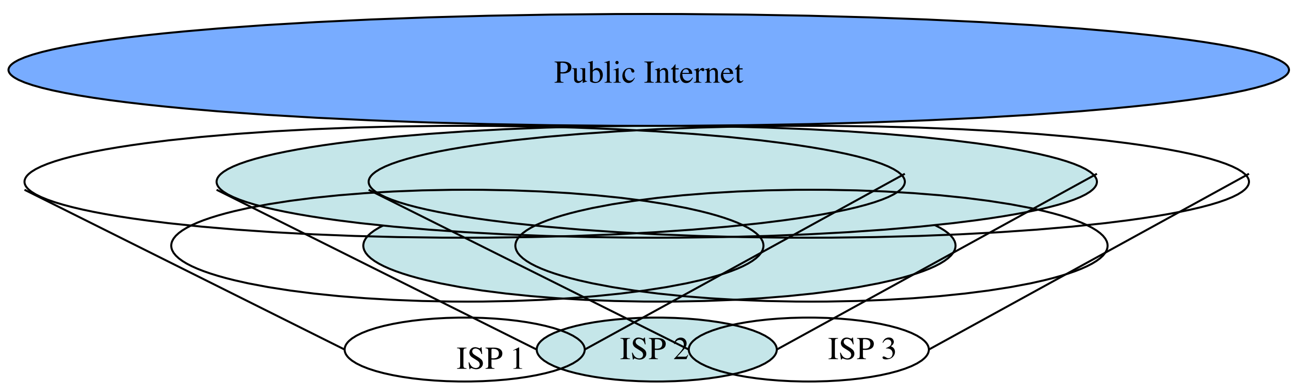

How Does It Work?

The Internet and ISPs

- But there does not need to be ONE e-mall.
  - You mean!
    - Yes, it is really an INTERnet!

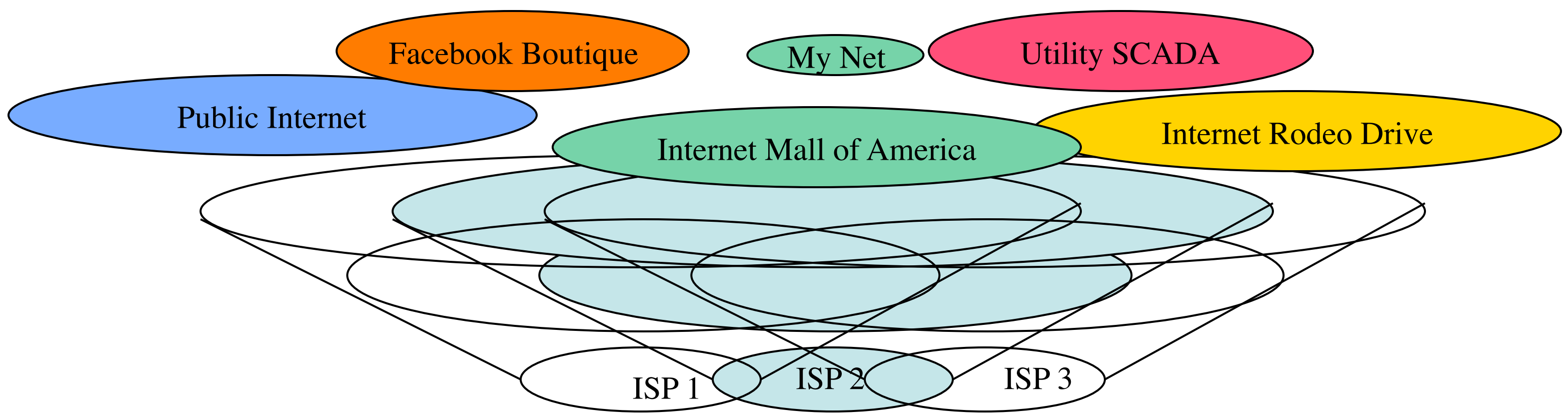

How Does It Work?

The User's Perspective

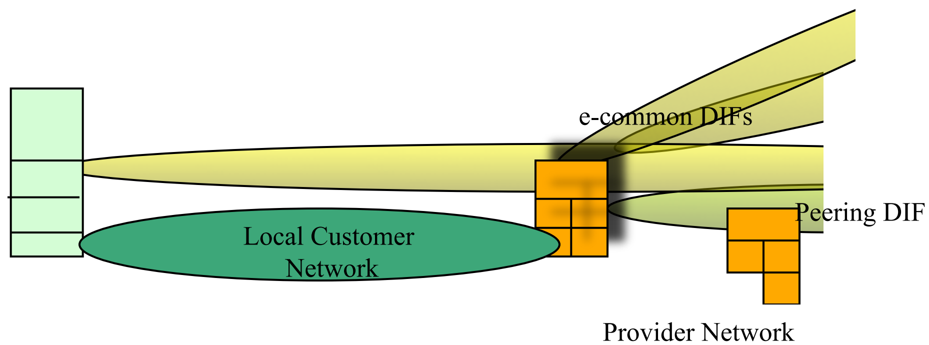

A Customer Network has a border router that makes several e-malls available. A choice can be made whether the entire local network joins, a single host or a single application.

In this case, one host on the local network chooses to join one of the available e-malls.

IRATI INVESTIGATING RINA

The Pouzin Society

Before Tackling Security
A Word on Method
(hardly news by now)

- When trying to work out the IPC Model absolutely no thought was given to security. All of the focus was just understanding the structure.
- People kept asking, What about Security? Is there a security layer?
- Didn't Know. Hadn't thought about it.
- There was the obvious:
  - The recursion of the layer provided Isolation.
  - That only the Application Name and local port-id were exposed to the correspondents.
- Interesting, but hardly an answer
- But it wasn't the time for those questions . . .
- At least not yet . . .

The Recursion Provided Isolation

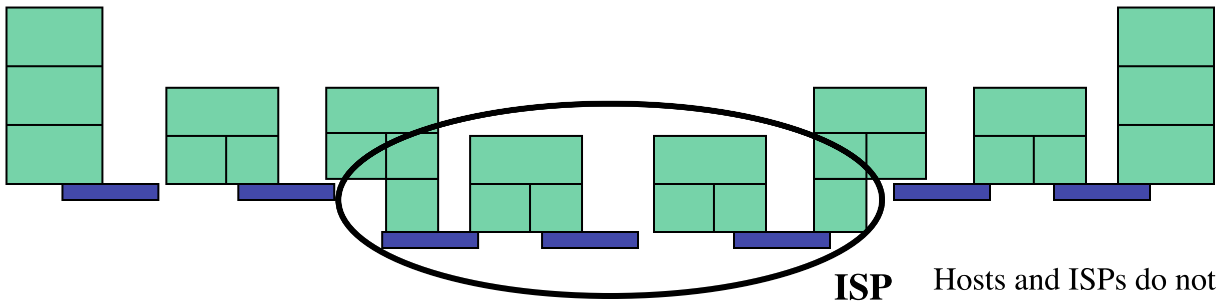

Hosts and ISPs do not share DIFS.
(ISP may have more layers

- Security by isolation, (not obscurity)
- Hosts can not address any element of the ISP.
- No user hacker can compromise ISP assets.
    - Unless ISP is physically compromised.

How Does It Work?

Security

- A Hacker in the Public Internet cannot connect to an Application in another DIF without either joining the DIF, or creating a new DIF spanning both. Either requires authentication and access control.
  – Non-IPC applications that can access two DIFs are a potential security problem.
- Certainly promising

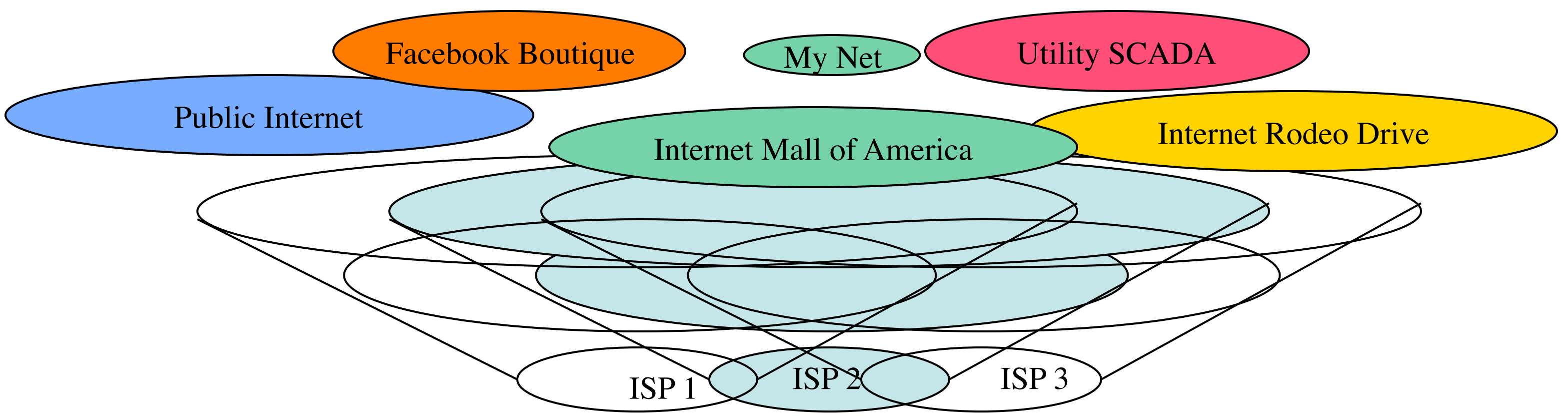

But When It Was Time

- The question was not, how to put in security?
- The question was,
- What does the IPC Model *tell* us about security?
  - Remember, our first task is always *understanding*.

- Let the Problem Answer the Question!
  - Let the Problem Tell Us What to Do.

The Problem Had a Lot to Say

The Pouzin Society

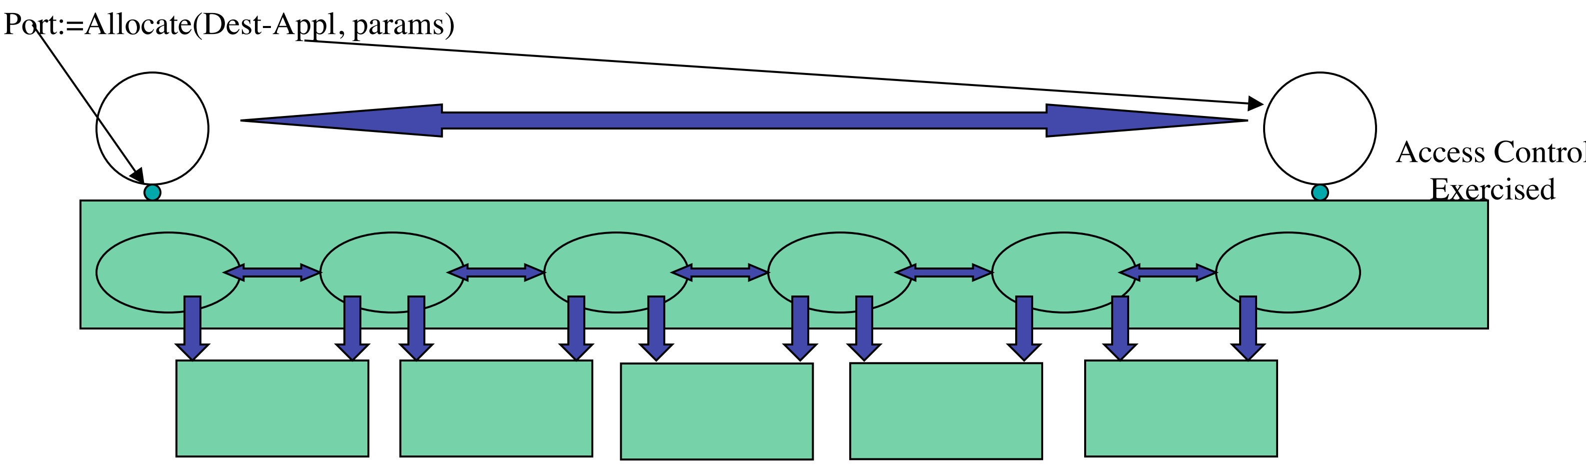

- We Already Mentioned How Little is Exposed the Layer Above.
- The Original OS Model indicated where Access Control went.
- Creating the Application Connection for Enrollment indicated where Authentication belonged, and that
  - Authentication of Applications must be done by the Applications themselves.
  - All members of the layer are authenticated within policy.
- SDU Protection clearly provided Confidentiality and Integrity.
- That implied that only Minimal trust was necessary:
  - Only that the lower layer will deliver something to someone.

A Very Unexpected Result

- A DIF with no explicit security mechanisms is inherently more secure than the current Internet under the same conditions!
- It would appear that
  - A DIF is a Securable Container.

Other Things Fall Into Place

- Data Transfer in RINA is based on Delta-t (Watson, 1980)
- Lot has happened in 30 years, many attacks on TCP have been found:
  - Port scanning
  - SYN attacks
  - Reset Attacks
  - Reassembly Attacks
- Long after delta-t was designed, what about delta-t?
- Short answer:
  - None of them work (Boddapati, et al., 2012)
    - Amazing, totally unexpected
  - Why not?
- Multiple fundamental reasons, but all inherent in the structure:
  - First, have to join the DIF (all members are authenticated)
  - Second, No Well-Known Ports
    - Would have to scan all possible application names!
  - Third and more importantly, . . .

Decoupling Port Allocation and Synchronization

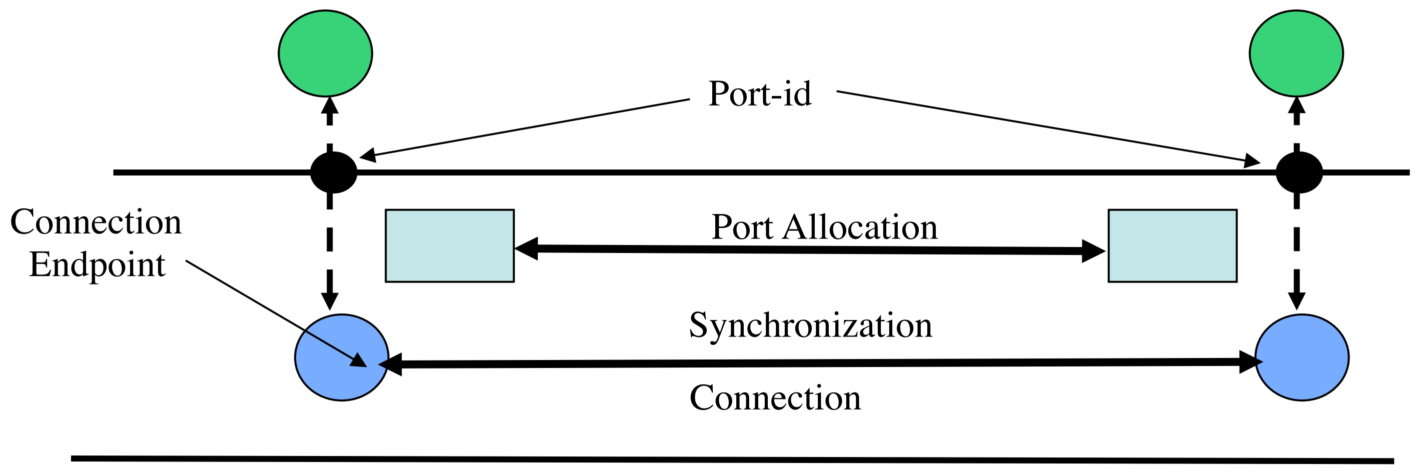

- No Way to Know What CEP-ids are Being Used, Since There is No Relation Between Port-id and CEP-id.
  - SYN-Ack Attack: must guess which of 2^16 CEP-id.
  - Data Transfer: must guess CEP-id and seq num within window!
  - Reassembly attack: Reassembly only done once.
  - No well-known ports to scan.

Decoupling Port Allocation and Synchronization: No IPSec

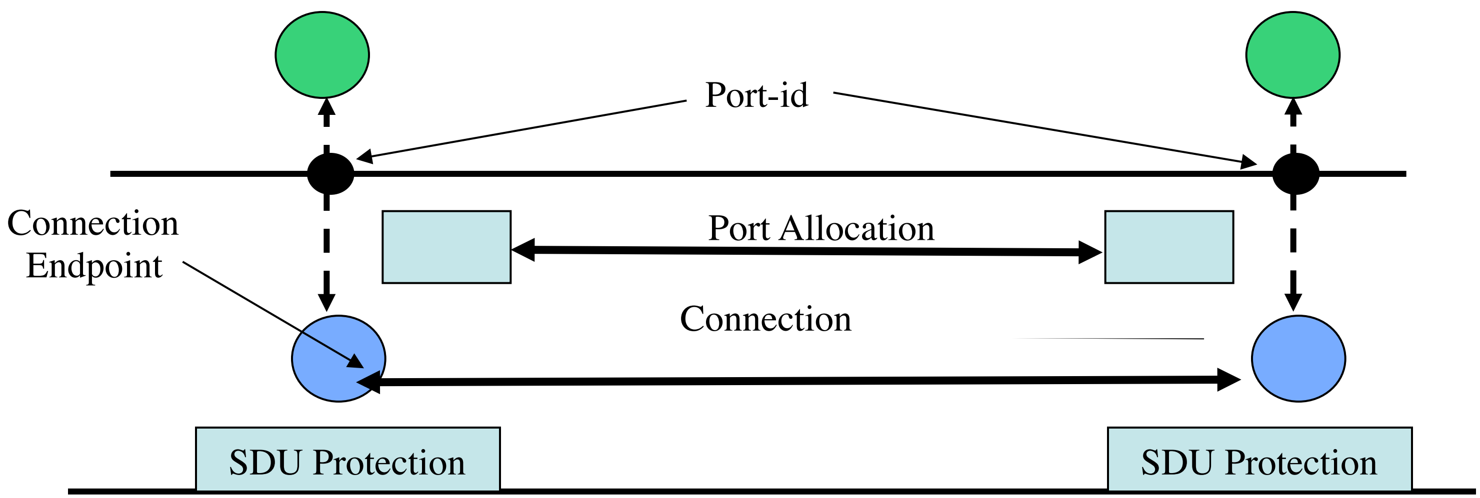

- IPsec is necessary with TCP/IP because no authentication and Sequence numbers turn over too quickly: don't repeat sequence number with same CEP-id.
- With RINA and delta-t, IPC Processes all authenticated, SDU Protection does the encryption, and packet sequence numbers slows rollover, but if it does, then simply allocate a new connection
- And bind it to the same port-ids, old one disappears after 2MPL.

RINA is Inherently More Secure and Less Work

The Pouzin Society

- A DIF is a Securable Container. (Small, 2011)
  - What info required to mount an attack, How to get the info
  - Small does a threat analysis at the architecture level
- Implies that Firewalls are Unnecessary,
  - The DIF *is* the Firewall!
- RINA Security is considerably Less Complex than the Current Internet Security (Small, 2012)
  - Only do a rough estimate counting protocols and mechanisms.
    - See paper for details.

To Add Security	Internet	RINA
Protocols	8	0
Non-Security Mechanisms	59	0
Security Mechanisms	28	7

Why Is Internet Security So Bad?

- The Standard Rationale One Sees is that They Didn't Think About It at the Beginning.
  - Neither did We.
  - Nor did Watson.
  - But RINA and delta-t are more secure.
- That Seems to Imply that
  - Good Design May be More Important to Security than Security Is.

Summing Up Security

- This is a MAJOR Improvement in Internet Security.
  - Not only more secure, but for less cost, with less overhead.
- So is Internet Security solved?
  - Hardly.
  - Still need: to develop the plug-in policy modules
  - to consider DDoS (we have some ideas)
  - As well as protecting against Rogue IPC Processes
  - and much more to explore and who knows what general principles will fall out.
- Most attacks are in the Applications, this does nothing about that.
  - But Much of this applies equally well to DAFs
    - Model implies that OS security reduces to Bounds Checking on Memory and IPC Security.
  - May also make it harder, might be able to deflect more DDoS attacks

There's More to Come

- Next Naming and Addressing
  - It turns out to be quite straightforward and simple.
- Then a Look at the Internal Operation of a DIF and the Implementations.
- A Claim: One will not find a structure that is both as rich and as simple as this that is not equivalent to it. Prove me wrong! ;-)
- But for Now. . .

IRATI
INVESTIGATING RINA

Pristine

The Pouzin Society

Questions?